

A Roadmap for the Computation of Persistent Homology

Nina Otter^{†§} Mason A. Porter^{†‡¶} Ulrike Tillmann^{†§} Peter Grindrod[†]
 Heather A. Harrington[†]

Abstract

Persistent homology (PH) is a method used in topological data analysis (TDA) to study qualitative features of data that persist across multiple scales. It is robust to perturbations of input data, independent of dimensions and coordinates, and provides a compact representation of the qualitative features of the input. There has been recent progress, but the computation of PH remains an open area with numerous important and fascinating challenges. The field of PH computation is evolving rapidly, and new algorithms and software implementations are being updated and released at a rapid pace. The purposes of our article are to (1) introduce theory and computational methods for PH to a broad range of computational scientists and (2) provide benchmarks of state-of-the-art implementations for the computation of PH. We give a friendly introduction to PH, navigate the pipeline for the computation of PH with an eye towards applications, and use a range of synthetic and real-world data sets to evaluate currently available open-source implementations for the computation of PH. Based on our benchmarking, we indicate which algorithms and implementations are best suited to different types of data sets. In an accompanying tutorial, we provide guidelines for the computation of PH. We make publicly available all scripts that we wrote for the tutorial, and we make available the processed version of the data sets used in the benchmarking.

1 Introduction

The amount of available data has increased dramatically in recent years, and this situation — which will only become more extreme — necessitates the development of innovative and efficient data-processing methods. Making sense of the vast amount of data is difficult: on one hand, the sheer size of the data poses challenges; on the other hand, the complexity of the data, which includes situations in which data is noisy, high-dimensional, and/or incomplete, is perhaps an even more significant challenge. The use of clustering techniques and other ideas from areas such as computer science, machine learning, and uncertainty quantification — along with mathematical and statistical models — are often very useful for data analysis (see, e.g., [60, 64, 78, 112] and many other references). However, recent mathematical developments are shedding new light on such “traditional” ideas, forging new approaches of their own, and helping people to better decipher increasingly complicated structure in data.

Techniques from the relatively new discipline of “topological data analysis” (TDA) have provided a wealth of new insights in the study of data in an increasingly diverse set of applications — including coverage in sensor networks [39], protein structure [59, 80, 130], stability of

²Mathematical Institute, University of Oxford, Oxford OX2 6GG, UK

³CABDyN Complexity Centre, University of Oxford, Oxford OX1 1HP, UK

⁴The Alan Turing Institute, 96 Euston Road, London NW1 2DB, UK

⁵Department of Mathematics, UCLA, Los Angeles, CA 90095, USA

fullerene molecules [129], robotics [11, 109, 123], signals in images [30, 68], periodicity in time series [106], breast-cancer classification [43, 103], viral evolution [26], natural images [24], the spread of contagions [90, 119], the structure of amorphous materials [70], force networks in granular matter [81, 82], equities-market networks [87], diverse applications in neuroscience [35, 63], collective behavior in biology [121], and time-series output of dynamical systems [91]. There are numerous others, and new applications of TDA appear in journals and preprint servers increasingly frequently. There are also interesting computational efforts, such as [127].

TDA is a field that lies at the intersection of data analysis, algebraic topology, computational geometry, computer science, statistics, and other related areas. The main goal of TDA is to use ideas and results from geometry and topology to develop tools for studying qualitative features of data. To achieve this goal, one needs precise definitions of qualitative features, tools to compute them in practice, and some guarantee about the robustness of those features. One way to address all three points is a method in TDA called *persistent homology* (PH). This method is appealing for applications because it is based on algebraic topology, which gives a well-understood theoretical framework to study qualitative features of data with complex structure, is computable via linear algebra, and is robust with respect to small perturbations in input data.

Types of data sets that can be studied with PH include finite metric spaces, digital images, level-sets of real-valued functions, and networks (see Section 5.1). In the next two paragraphs, we give some motivation for the main ideas of persistent homology by discussing two examples of such data sets.

Finite metric spaces are also called “point cloud” data sets in the TDA literature. From a topological point of view, finite metric spaces do not contain any interesting information. One thus considers a “thickening” of a point cloud at different “scales of resolution” and then analyses the evolution of the resulting shape across the different resolution scales. The qualitative features are given by topological invariants, and one can represent the variation of such invariants across the different resolution scales in a compact way to summarize the “shape” of the data.

As an illustration, consider the set of points in \mathbb{R}^2 that we show in Fig. 1. Let ϵ , which we interpret as a “distance parameter,” be a nonnegative real number (so $\epsilon = 0$ gives the set of points). For different values of ϵ , we construct a space S_ϵ composed of vertices, edges, triangles, and higher-dimensional polytopes according to the following rule: We include an edge between two points i and j if and only if the Euclidean distance between them is no larger than ϵ ; we include a triangle if and only if all of its edges are in S_ϵ ; we include a tetrahedron if and only if all of its face triangles are in S_ϵ ; and so on. For $\epsilon \leq \epsilon'$, it then follows that the space S_ϵ is contained in the space $S_{\epsilon'}$. This yields a nested sequence of spaces, as we illustrate in Fig. 1(a).

By using homology, a tool in algebraic topology, one can measure several features of the spaces S_ϵ — including the numbers of components, holes, and voids (higher-dimensional versions of holes). One can then represent the lifetime of such features using a finite collection of intervals known as a “barcode.” Roughly, the left endpoint of an interval represents the birth of a feature, and its right endpoint represents the death of the same feature. In Fig. 1(b), we reproduce such intervals for the number of components (blue solid lines) and the number of holes (violet dashed lines).

In Fig. 1(b), we observe a dashed line that is significantly longer than the other dashed lines. This indicates that the data set has a long-lived hole. By contrast, one can potentially construe the shorter dashed lines as noise¹. When a feature is still “alive” at the largest value of ϵ that we consider, the lifetime interval is an infinite interval, which we indicate by putting an arrowhead

¹However, note that shorter lines cannot always be construed as noise (see, e.g., [117]).




Figure 1: (a) A finite set of points in \mathbb{R}^2 (for $\epsilon = 0$) and a nested sequence of spaces obtained from it (from $\epsilon = 0$ to $\epsilon = 2.1$). (b) Barcode for the nested sequence of spaces illustrated in (a). Solid lines represent the lifetime of components, and dashed lines represent the lifetime of holes.

at the right endpoint of the interval. In Fig. 1(b), we see that there is exactly one solid line that lives up to $\epsilon = 2.1$. One can use information about shorter solid lines to extract information about how data is clustered in a similar way as with linkage-clustering methods [60]. Note that one of the most difficult parts of using PH is the statistical interpretation of results. We discuss such interpretation further in Section 5.4. Our construction of nested spaces gives an example of a “filtered Vietoris–Rips complex,” which we define and discuss in Section 5.2.

We now discuss a second example related to digital images. (For an illustration, see Fig. 2(a).) Digital images have a “cubical” structure, given by the pixels (for 2-dimensional digital images) or voxels (for 3-dimensional images). Therefore, one approach to study digital images uses combinatorial structures called “cubical complexes.” (For a different approach to the study of digital images, see Section 5.1.) Roughly, cubical complexes are topological spaces built from a union of vertices, edges, squares, cubes, and higher-dimensional hypercubes. An efficient way [126] to build a cubical complex from a 2-dimensional digital image consists of assigning a vertex to every pixel, then joining vertices corresponding to adjacent pixels by an edge, and filling in the resulting squares. . . One proceeds in a similar way for 3-dimensional images. One then labels every vertex with an integer that corresponds to the gray value of the pixel, and one labels edges (respectively, squares) with the maximum of the values of the adjacent vertices (respectively, edges). One can then construct a nested sequence of cubical complexes $C_0 \subset C_1 \subset \dots \subset C_{256}$, where for each $i \in \{0, 1, \dots, 256\}$, the cubical complex C_i contains all vertices, edges, squares, and cubes that are labeled by a number less than or equal to i . (See Fig. 2(c) for an example.) Such a sequence of cubical complexes is also called a “filtered cubical complex.” Similar to the previous example, one can use homology to measure several features of the spaces C_i (see Fig. 2(d)).

In the present article, we focus on persistent homology, but there are also other methods in TDA — including the Mapper algorithm [114], Euler calculus (see [62] for an introduction with




Figure 2: (a) A gray-scale image, (b) the matrix of gray values, (c) the filtered cubical complex associated to the digital image, and (d) the barcode for the nested sequence of spaces in panel (c). A solid line represents the lifetime of a component, and a dashed line represents the lifetime of a hole.

an eye towards applications), cellular sheaves [34,62], and many more. We refer readers who wish to learn more about the foundations of TDA to the article [22], which discusses why topology and functoriality are essential for data analysis. We point to several introductory papers, books, and two videos on PH at the end of Section 4.

The first algorithm for the computation of PH was introduced for computation over \mathbb{F}_2 (the field with two elements) in [53] and over general fields in [137]. Since then, several algorithms and optimization techniques have been presented, and there are now various powerful implementations of PH [5,6,9,93,97,100,118]. Those wishing to try PH for computations may find it difficult to discern which implementations and algorithms are best suited for a given task. The field of PH is evolving continually, and new software implementations and updates are released at a rapid pace. Not all of them are well-documented, and (as is well-known in the TDA community), the computation of PH for large data sets is computationally very expensive.

To our knowledge, there exists neither an overview of the various computational methods for PH nor a comprehensive benchmarking of the state-of-the-art implementations for the computation of persistent homology. In the present article, we close this gap: we introduce computation of PH to a general audience of applied mathematicians and computational scientists, offer guidelines for the computation of PH, and test the existing open-source published libraries for the computation of PH.

The rest of our paper is organized as follows. In Section 2, we discuss related work. We then introduce homology in Section 3 and introduce PH in Section 4. We discuss the various steps of the pipeline for the computation of PH in Section 5, and we briefly examine algorithms for “generalized persistence” in Section 6. In Section 7, we give an overview of software libraries, discuss our benchmarking of a collection of them, and provide guidelines for which software or algorithm is better suited to which data set. (We provide specific guidelines for the computation of PH with the different libraries in the Tutorial in the Supplementary Information (SI).) In Section 8, we discuss future directions for the computation of PH.

2 Related work

In our work, we introduce PH to non-experts with an eye towards applications, and we benchmark state-of-the-art libraries for the computation of PH. In this section, we discuss related work for both of these points.

There are several excellent introductions to the theory of PH (see the references at the end of Section 4.1), but none of them emphasizes the actual computation of PH by providing specific guidelines for people who want to do computations. In the present paper, we navigate the theory of PH with an eye towards applications, and we provide guidelines for the computation of PH using the open-source libraries PERSEUS, DIONYSUS, DIPHA, JAVAPLEX, and GUDHI. We include a tutorial (see the SI) that gives specific instructions for how to use the different functionalities that are implemented in these libraries. Much of this information is scattered throughout numerous different papers, websites, and even source code of implementations, and we believe that it is beneficial to the applied mathematics community (especially people who seek an entry point into PH) to find all of this information in one place. The functionalities that we cover include plots of barcodes and persistence diagrams and the computation of PH with Vietoris–Rips complexes, alpha complexes, Čech complexes, witness complexes, and cubical complexes for image data. We thus believe that our paper closes a gap in introducing PH to people interested in applications, while our tutorial complements existing tutorials (see, e.g. [2,19,57]).

We believe that there is a need for a thorough benchmarking of the state-of-the-art libraries. In our work, we use twelve different data sets to test and compare the libraries JAVAPLEX,

GUDHI, DIPHA, DIONYSUS, and PERSEUS, and we obtain some surprising results (see Section 7.2). There are several benchmarkings in the PH literature; we are aware of the following ones: the benchmarking in [40] compares the implementations of standard and dual algorithms in DIONYSUS; the one in [101] compares the Morse-theoretic reduction algorithm with the standard algorithm; the one in [9] compares all of the data structures and algorithms implemented in PHAT; the benchmarking in [8] compares PHAT and its spin-off DIPHA; and the benchmarking in C. Maria’s doctoral thesis [92] is to our knowledge the only existing benchmarking that compares packages from different authors. However, Maria compares only up to three different implementations at one time, and he used the package JPLEX (which is no longer maintained) instead of the JAVAPLEX library (its successor). Additionally, the widely used library PERSEUS (e.g., it was used in [81, 82, 115, 119]) does not appear in Maria’s benchmarking.

3 Homology

Assume that one is given data that lies in a metric space, such as a subset of Euclidean space with an inherited distance function. In many situations, one is not interested in the precise geometry of these spaces, but instead seeks to understand some basic characteristics, such as the number of components or the existence of holes and voids. Algebraic topology captures these basic characteristics either by counting them or by associating vector spaces or more sophisticated algebraic structures to them. Here we are interested in *homology*, which associates one vector space $H_i(X)$ to a space X for each natural number $i \in \{0, 1, 2, \dots\}$. The dimension of $H_0(X)$ counts the number of path components in X , the dimension of $H_1(X)$ is a count of the number of holes, and the dimension of $H_2(X)$ is a count of the number of voids. An important property of these algebraic structures is that they are robust, as they do not change when the underlying space is transformed by bending, stretching, and other deformations. In technical terms, they are *homotopy invariant*. Conversely, under favorable conditions², these algebraic invariants determine the topology of a space up to homotopy — an equivalence relation that is much coarser (and easier to work with) than the more familiar notion of homeomorphy.

It can be very difficult to compute the homology of arbitrary topological spaces. We thus approximate our spaces by combinatorial structures called “simplicial complexes,” for which homology can be easily computed algorithmically. Indeed, often one is not even given the space X , but instead possesses only a discrete sample set S from which to build a simplicial complex following one of the recipes described in Sections 3.2 and 5.2.

3.1 Simplicial complexes and their homology

We begin by giving the definitions of simplicial complexes and of the maps between them. Roughly, a simplicial complex is a space that is built from a union of points, edges, triangles, tetrahedra, and higher-dimensional polytopes. We illustrate the main definitions given in this section with the example in Fig. 3. As we pointed out in Section 1, “cubical complexes” give another way to associate a combinatorial structure to a topological space. In TDA, cubical complexes have been used primarily to study image data sets. One can compute PH for a nested sequence of cubical complexes in a similar way as for simplicial complexes, but the theory of PH for simplicial complexes is richer, and we therefore examine only simplicial homology and complexes in our discussions.

²See [69, Corollary 4.33].

Definition 1 A simplicial complex³ is a collection K of non-empty subsets of a set K_0 such that $\tau \subset \sigma$ and $\sigma \in K$ guarantees that $\tau \in K$ and $\{v\} \in K$ for all $v \in K_0$. The elements of K_0 are called vertices of K , and the elements of K are called simplices. Additionally, we say that a simplex has dimension p or is a p -simplex if it has a cardinality of $p + 1$. We use K_p to denote the collection of p -simplices. The k -skeleton of K is the union of the sets K_p for all $p \in \{0, 1, \dots, k\}$. If τ and σ are simplices such that $\tau \subset \sigma$, then we call τ a face of σ . The dimension of K is defined as the maximum of the dimensions of its simplices. A map of simplicial complexes, $f: K \rightarrow L$, is a map $f: K_0 \rightarrow L_0$ such that $f(\sigma) \in L$ for all $\sigma \in K$.

We give an example of a simplicial complex in Fig. 3(a) and an example of a map of simplicial complexes in Fig. 3(b). Definition 1 is rather abstract, but one can always interpret a finite simplicial complex K geometrically as a subset of \mathbb{R}^N for sufficiently large N ; such a subset is called a “geometric realization,” and it is unique up to a canonical piecewise-linear homeomorphism. For example, the simplicial complex in Fig. 3(a) has a geometric realization given by the subset of \mathbb{R}^2 in Fig. 3(c).

(a) A simplicial complex:

$$K = \{\{a\}, \{b\}, \{c\}, \{d\}, \{e\}, \{a, b\}, \{a, c\}, \{a, d\}, \{b, c\}, \{c, d\}, \{a, b, c\}\}.$$

(b) Let K be the simplicial complex in (a), and let K' be the following simplicial complex:

$$K' = \{\{g\}, \{h\}, \{l\}, \{g, h\}, \{g, l\}, \{h, l\}, \{g, h, l\}\}.$$

We define a map of simplicial complexes, $f: K \rightarrow K'$, as follows:

$$f: K_0 \rightarrow K'_0: x \mapsto \begin{cases} h, & \text{if } x=a \\ g, & \text{otherwise.} \end{cases}$$

(c) Geometric realizations of the simplicial complexes in (a) and (b) are the following subsets of \mathbb{R}^2 :



Figure 3: (a) Example of a simplicial complex, (b) a map of simplicial complexes, and (c) a geometric realization of the simplicial complex in (a).

We now define homology for simplicial complexes. Let \mathbb{F}_2 denote the field with two elements. Given a simplicial complex K , let $C_p(K)$ denote the \mathbb{F}_2 -vector space with basis given by the p -simplices of K . For any $p \in \{1, 2, \dots\}$, we define the linear map

$$d_p: C_p(K) \rightarrow C_{p-1}(K): \sigma \mapsto \sum_{\tau \subset \sigma, \tau \in K_{p-1}} \tau.$$

For $p = 0$, we define d_0 to be the zero map. In words, d_p maps each p -simplex to its boundary, the sum of its faces of codimension 1. Because the boundary of a boundary is always empty, the

³Note that this is usually called an “abstract simplicial complex” in the literature.

linear maps d_p have the property that composing any two consecutive maps yields the zero map: for all $p \in \{0, 1, 2, \dots\}$, we have $d_p \circ d_{p+1} = 0$. Consequently, the image of d_{p+1} is contained in the kernel of d_p , so we can take the quotient of $\text{kernel}(d_p)$ by $\text{image}(d_{p+1})$. We can thus make the following definition.

Definition 2 For any $p \in \{0, 1, 2, \dots\}$, the p th homology of a simplicial complex K is the quotient vector space

$$H_p(K) := \text{kernel}(d_p) / \text{image}(d_{p+1}).$$

Its dimension

$$\beta_p(K) := \dim H_p(K) = \dim \text{kernel}(d_p) - \dim \text{image}(d_{p+1})$$

is called the p th Betti number of K . Elements in the image of d_{p+1} are called p -boundaries, and elements in the kernel of d_p are called p -cycles.

Intuitively, the p -cycles that are not boundaries represent p -dimensional holes. Therefore, the p th Betti number “counts” the number of p -holes. Additionally, if K is a simplicial complex of dimension n , then for all $p > n$, we have that $H_p(K) = 0$, as K_p is empty and hence $C_p(K) = 0$. We therefore obtain the following sequence of vector spaces and linear maps:

$$0 \xrightarrow{d_{n+1}} C_n(K) \xrightarrow{d_n} \dots \xrightarrow{d_2} C_1(K) \xrightarrow{d_1} C_0(K) \xrightarrow{d_0} 0.$$

We give an example of such a sequence in Fig. 4(a), for which we also report the Betti numbers.

One of the most important properties of simplicial homology is “functoriality.” Any map $f: K \rightarrow K'$ of simplicial complexes induces the following \mathbb{F}_2 -linear map:

$$\tilde{f}_p: C_p(K) \rightarrow C_p(K'): \sum_{\sigma \in K_p} c_\sigma \sigma \mapsto \sum_{\sigma \in K_p \text{ such that } f(\sigma) \in K'_p} c_\sigma f(\sigma), \quad \text{for any } p \in \{0, 1, 2, \dots\},$$

where $c_\sigma \in \mathbb{F}_2$. Additionally, $\tilde{f}_p \circ d_{p+1} = d'_{p+1} \circ \tilde{f}_{p+1}$, and the map \tilde{f}_p therefore induces the following linear map between homology vector spaces:

$$\begin{aligned} f_p: H_p(K) &\rightarrow H_p(K'), \\ [c] &\mapsto [\tilde{f}_p(c)]. \end{aligned}$$

(We give an example of such a map in Fig. 4(b).) Consequently, to any map $f: K \rightarrow K'$ of simplicial complexes, we can assign a map $f_p: H_p(K) \rightarrow H_p(K')$ for any $p \in \{0, 1, 2, \dots\}$. This assignment has the important property that given a pair of composable maps of simplicial complexes, $f: K \rightarrow K'$ and $g: K' \rightarrow K''$, the map $(g \circ f)_p: H_p(K) \rightarrow H_p(K'')$ is equal to the composition of the maps induced by f and g . That is, $(g \circ f)_p = g_p \circ f_p$. The fact that a map of simplicial complexes induces a map on homology that is compatible with composition is called *functoriality*, and it is crucial for the definition of persistent homology (see Section 4.1).

When working with simplicial complexes, one can modify a simplicial complex by removing or adding a pair of simplices (σ, τ) , where τ is a codimension-1 face of σ and σ is the only simplex that has τ as a face. The resulting simplicial complex has the same homology as the one with which we started. In Fig. 3 (a), we can remove the pair $(\{a, b, c\}, \{b, c\})$ and then the pair $(\{a, b\}, \{b\})$ without changing the Betti numbers. Such a move is called an *elementary simplicial collapse* [31]. In Section 5.2.6, we will see an application of this for the computation of PH.

In this section, we have defined simplicial homology over the field \mathbb{F}_2 — i.e., “with coefficients in \mathbb{F}_2 .” One can be more general and instead define simplicial homology with coefficients in any field (or even in the integers). However, when $1 \neq -1$, one needs to take more care when defining

- (a) We compute the simplicial homology for the simplicial complex in Fig. 3(a). We have the following sequence of vector spaces and linear maps:

$$0 \longrightarrow \mathbb{F}_2 \xrightarrow{d_2} \mathbb{F}_2^5 \xrightarrow{d_1} \mathbb{F}_2^5 \xrightarrow{d_0} 0.$$

Let abc denote the basis vector that corresponds to the simplex $\{a, b, c\}$. Similarly, we use ab , ac , ad , bc , and cd to denote the basis vectors that correspond to the 1-simplices; and we use a , b , c , d , and e to denote the basis vectors that correspond to the 0-simplices. We order the bases of the vector spaces using lexicographic order. We then have

$$d_2 = (1 \quad 1 \quad 0 \quad 1 \quad 0)^t$$

and

$$d_1 = \begin{pmatrix} 1 & 1 & 1 & 0 & 0 \\ 1 & 0 & 0 & 1 & 0 \\ 0 & 1 & 0 & 1 & 1 \\ 0 & 0 & 1 & 0 & 1 \\ 0 & 0 & 0 & 0 & 0 \end{pmatrix}.$$

One can then compute that $\beta_0(K) = 2$, $\beta_1(K) = 1$, and all higher Betti numbers are 0.

- (b) The map of simplicial complexes, $f: K \rightarrow K'$, in Fig. 3(b) induces the map $\tilde{f}_0: C_0(K) \rightarrow C_0(K')$ on chain complexes that sends the basis element a to h (note that we use the same notation as in part (a) of this box). Furthermore, it sends the basis elements that correspond to the other 0-simplices of K to g , and it induces the map $\tilde{f}_1: C_1(K) \rightarrow C_1(K')$ that sends the basis elements ab , ac , and ad to cd and sends the basis elements cd and bc to 0. This gives a map f_0 on degree-0 homology that sends both generators of $H_0(K)$ to the generator of $H_0(K')$. It follows that $H_i(K') = 0$ for all $i \geq 1$, so f_i is the zero map for $i \geq 1$.

Figure 4: (a) Computation of simplicial homology for the simplicial complex in Fig. 3(a) and (b) induced map in degree-0 homology for the map of simplicial complexes in Fig. 3(b).

the boundary maps d_p to ensure that $d_p \circ d_{p+1}$ remains the zero map. Consequently, the definition is more involved. For the purposes of the present paper, it suffices to consider homology with coefficients in the field \mathbb{F}_2 . Indeed, we will see in Section 4 that to obtain topological summaries in the form of barcodes, we need to compute homology with coefficients in a field. Furthermore, as we summarize in Table 2 (in Section 7), most of the implementations for the computation of PH work with \mathbb{F}_2 . See [69] for an introduction to homology, and see [76] for an overview of computational homology.

3.2 Building simplicial complexes

As we discussed in Section 3.1, computing the homology of finite simplicial complexes boils down to linear algebra. The same is not true for the homology of an arbitrary space X , and one therefore tries to find simplicial complexes whose homology approximates the homology of the

space in an appropriate sense.

An important tool is the Čech (\check{C}) complex. Let \mathcal{U} be a cover of X — i.e., a collection of subsets of X such that the union of the subsets is X . The k -simplices of the Čech complex are the non-empty intersections of $k + 1$ sets in the cover \mathcal{U} . More precisely, we define the nerve of a collection of sets as follows.

Definition 3 Let $\mathcal{U} = \{U_i\}_{i \in I}$ be a non-empty collection of sets. The nerve of \mathcal{U} is the simplicial complex with set of vertices given by I and k -simplices given by $\{i_0, \dots, i_k\}$ if and only if $U_{i_0} \cap \dots \cap U_{i_k} \neq \emptyset$.

If the cover of the sets is sufficiently “nice,” then the Nerve Theorem implies that the nerve of the cover and the space X have the same homology [13, 51]. For example, suppose that we have a finite set of points S in a metric space X . We then can define, for every $\epsilon > 0$, the space S_ϵ as the union $\cup_{x \in S} B(x, \epsilon)$, where $B(x, \epsilon)$ denotes the closed ball with radius ϵ centered at x . It follows that $\{B(x, \epsilon) \mid x \in S\}$ is a cover of S_ϵ , and the nerve of this cover is the Čech complex on S at scale ϵ . We denote this complex by $\check{C}_\epsilon(S)$. If the space X is Euclidean space, then the Nerve Theorem guarantees that the simplicial complex $\check{C}_\epsilon(S)$ recovers the homology of S_ϵ .

From a computational point of view, the Čech complex is expensive because one has to check for large numbers of intersections. Additionally, in the worst case, the Čech complex can have dimension $|\mathcal{U}| - 1$, and it therefore can have many simplices in dimensions higher than the dimension of the underlying space. Ideally, it is desirable to construct simplicial complexes that approximate the homology of a space but are easy to compute and have “few” simplices, especially in high dimensions. This is a subject of ongoing research: In Subsection 5.2, we give an overview of state-of-the-art methods to associate complexes to point-cloud data in a way that addresses one or both of these desiderata. See [51, 56] for more details on the Čech complex, and see [13, 51] for a precise statement of the Nerve Theorem.

4 Persistent homology

Assume that we are given experimental data in the form of a finite metric space S ; there are points or vectors that represent measurements along with some distance function (e.g., given by a correlation or a measure of dissimilarity) on the set of points or vectors. Whether or not the set S is a sample from some underlying topological space, it is useful to think of it in those terms. Our goal is to recover the properties of such an underlying space in a way that is robust to small perturbations in the data S . In a broad sense, this is the subject of *topological inference*. (See [104] for an overview.) If S is a subset of Euclidean space, one can consider a “thickening” S_ϵ of S given by the union of balls of a certain fixed radius ϵ around its points and then compute the Čech complex. One can thus try to compute qualitative features of the data set S by constructing the Čech complex for a chosen value ϵ and then computing its simplicial homology. The problem with this approach is that there is a priori no clear choice for the value of the parameter ϵ . The key insight of PH is the following: To extract qualitative information from data, one considers several (or even all) possible values of the parameter ϵ . As the value of ϵ increases, simplices are added to the complexes. Persistent homology then captures how the homology of the complexes changes as the parameter value increases, and it detects which features “persist” across changes in the parameter value. We give an example of persistent homology in Fig. 5.

4.1 Filtered complexes and homology

Let K be a finite simplicial complex, and let $K_1 \subset K_2 \subset \dots \subset K_l = K$ be a finite sequence of nested subcomplexes of K . The simplicial complex K with such a sequence of subcomplexes is

called a *filtered simplicial complex*. See Fig. 5(a) for an example of filtered simplicial complex. We can apply homology to each of the subcomplexes. For all homology degrees p , the inclusion maps $K_i \rightarrow K_j$ induce \mathbb{F}_2 -linear maps $f_{i,j}: H_p(K_i) \rightarrow H_p(K_j)$ for all $i, j \in \{1, \dots, l\}$ with $i \leq j$. By functoriality (see Section 3.1), it follows that

$$f_{k,j} \circ f_{i,k} = f_{i,j} \quad \text{for all } i \leq k \leq j. \quad (1)$$

We therefore give the following definition⁴.

Definition 4 Let $K_1 \subset K_2 \subset \dots \subset K_l = K$ be a filtered simplicial complex. The p th persistent homology of K is the pair

$$\left(\{H_p(K_i)\}_{1 \leq i \leq l}, \{f_{i,j}\}_{1 \leq i \leq j \leq l} \right),$$

where for all $i, j \in \{1, \dots, l\}$ with $i \leq j$, the linear maps $f_{i,j}: H_p(K_i) \rightarrow H_p(K_j)$ are the maps induced by the inclusion maps $K_i \rightarrow K_j$.

The p th persistent homology of a filtered simplicial complex gives more refined information than just the homology of the single subcomplexes. We can visualize the information given by the vector spaces $H_p(K_i)$ together with the linear maps $f_{i,j}$ by drawing the following diagram: at filtration step i , we draw as many bullets as the dimension of the vector space $H_p(K_i)$. We then connect the bullets as follows: we draw an interval between bullet u at filtration step i and bullet v at filtration step $i+1$ if the generator of $H_p(K_i)$ that corresponds to u is sent to the generator of $H_p(K_{i+1})$ that corresponds to v . If the generator corresponding to a bullet u at filtration step i is sent to 0 by $f_{i,i+1}$, we draw an interval starting at u and ending at $i+1$. (See Fig. 5(b) for an example.) Such a diagram clearly depends on a choice of basis for the vector spaces $H_p(K_i)$, and a poor choice can lead to complicated and unreadable clutter. Fortunately, by the Fundamental Theorem of Persistent Homology [137], there is a choice of basis vectors of $H_p(K_i)$ for each $i \in \{1, \dots, l\}$ such that one can construct the diagram as a well-defined and unique collection of disjoint half-open intervals, collectively called a *barcode*⁵. We give an example of a barcode in Fig. 5(c). Note that the Fundamental Theorem of PH, and hence the existence of a barcode, relies on the fact that we are using homology with field coefficients. (See [137] for more details.)

There is a useful interpretation of barcodes in terms of births and deaths of generators. Considering the maps $f_{i,j}$ written in the basis given by the Fundamental Theorem of Persistent Homology, we say that $x \in H_p(K_i)$ (with $x \neq 0$) is *born* in $H_p(K_i)$ if it is not in the image of $f_{i-1,i}$ (i.e., $f_{i-1,i}^{-1}(x) = \emptyset$). For $x \in H_p(K_i)$ (with $x \neq 0$), we say that x *dies* in $H_p(K_j)$ if $j > i$ is the smallest index for which $f_{i,j}(x) = 0$. The lifetime of x is represented by the half-open interval $[i, j)$. If $f_{i,j}(x) \neq 0$ for all $j > i$ in I , we say that x *lives forever*, and its lifetime is represented by the interval $[i, \infty)$.

Remark 5 Note that some references (e.g., [51]) introduce persistent homology by defining the birth and death of generators without using the existence of a choice of compatible bases, as given by the Fundamental Theorem of Persistent Homology. The definition of birth coincides with the

⁴A pair $(\{M_i\}_{i \in I}, \{\phi_{i,j}: M_i \rightarrow M_j\}_{i \leq j})$, where (I, \leq) is a totally ordered set, such that for each i , we have that M_i is a vector space and the maps $\phi_{i,j}$ are linear maps satisfying the functoriality property (1), is usually called a *persistence module*. With this terminology, the homology of a filtered simplicial complex is an example of persistence module.

⁵Although the collection of intervals is unique, note that one has to choose a vertical order when drawing the intervals in the diagram, and there is therefore an ambiguity in the representation of the intervals as a barcode. However, there is no ambiguity when representing the intervals as points in a persistence diagram (see Fig. 5(d)).

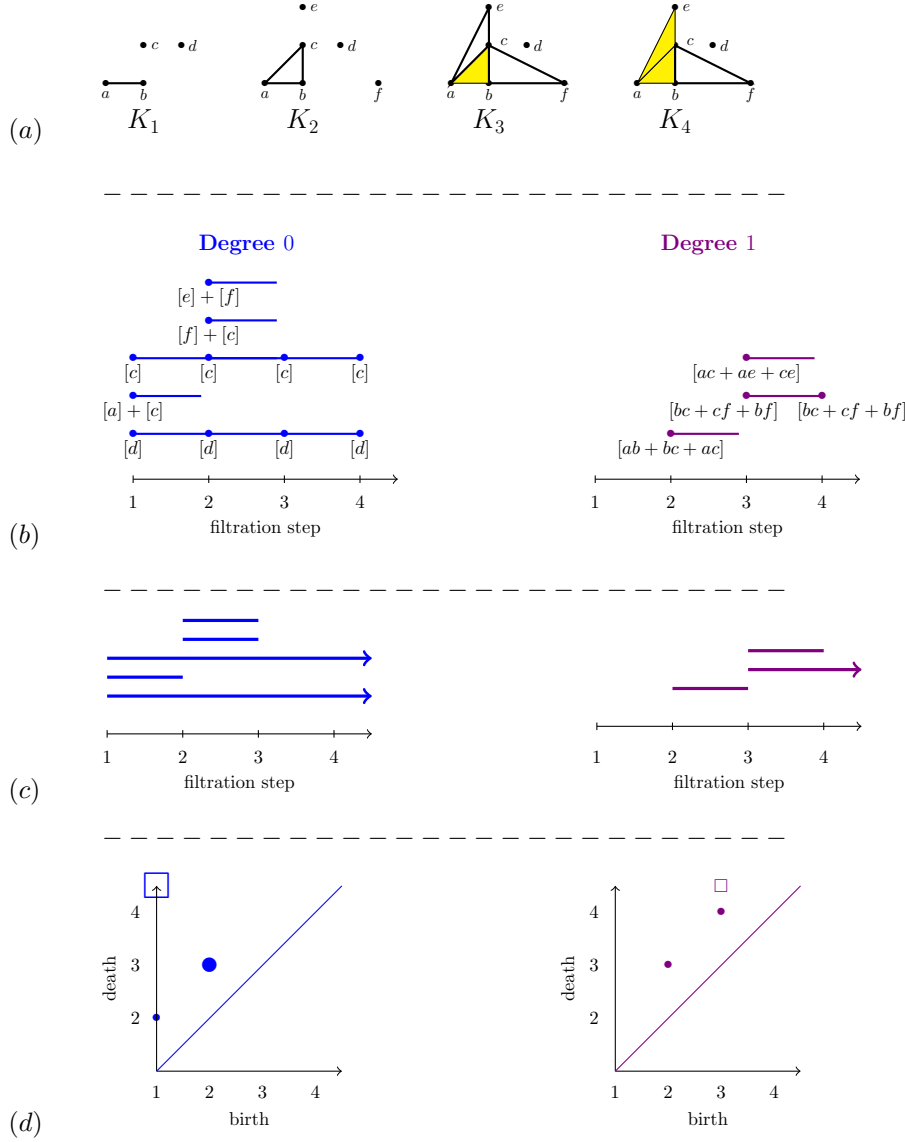


Figure 5: Example of persistent homology for a finite filtered simplicial complex. (a) We start with a finite filtered simplicial complex. (b) At each filtration step i , we draw as many vertices as the dimension of (left column) $H_0(K_i)$ and (right column) $H_1(K_i)$. We label the vertices by basis elements, the existence of which is guaranteed by the Fundamental Theorem of Persistent Homology, and we draw an edge between two vertices to represent the maps $f_{i,j}$, as explained in the main text. We thus obtain a well-defined collection of disjoint half-open intervals called a “barcode.” We interpret each interval in degree p as representing the lifetime of a p -homology class across the filtration. (c) We rewrite the diagram in (b) in the conventional way. We represent classes that are born but do not die at the final filtration step using arrows that start at the birth of that feature and point to the right. (d) An alternative graphical way to represent barcodes (which gives exactly the same information) is to use *persistence diagrams*, in which an interval $[i, j)$ is represented by the point (i, j) in the extended plane \mathbb{R}^2 , where $\mathbb{R} = \mathbb{R} \cup \{\infty\}$. Therefore, a persistence diagram is a finite multiset of points in \mathbb{R}^2 . We use squares to signify the classes that do not die at the final step of a filtration, and the size of dots or squares is directly proportional to the number of points being represented. For technical reasons, which we discuss briefly in Section 5.4, one also adds points on the diagonal to the persistence diagrams. (Each of the points on the diagonal has infinite multiplicity.)

definition that we have given, but the definition of death is different. One says that $x \in H_p(K_i)$ (with $x \neq 0$) dies in $H_p(K_j)$ if $j > i$ is the smallest index for which either $f_{i,j}(x) = 0$ or there exists $y \in H_p(K_{i'})$ with $i' < i$ such that $f_{i',j}(y) = f_{i,j}(x)$. In words, this means that x and y merge at filtration step j , and the class that was born earlier is the one that survives. In the literature, this is called the elder rule. We do not adopt this definition, because the elder rule is not well-defined when two classes are born at the same time, as there is no way to choose which class will survive. For example, in Fig. 5, there are two classes in H_0 that are born at the same stage in K_1 . These two classes merge in K_2 , but neither dies. The class that dies is $[a] + [c]$.

There are numerous excellent introductions to PH, such as the books [51,62,104,133] and the papers [22,49,50,61]. For a brief and friendly introduction to PH and some of its applications, see the video <https://www.youtube.com/watch?v=h0bnG1Wavag>. For a brief introduction to some of the ideas in TDA, see the video <https://www.youtube.com/watch?v=XfWibrh6stw>.

5 Computation of PH for data

We summarize the pipeline for the computation of PH from data in Fig. 6. In the following subsections, we describe each step of this pipeline and state-of-the-art algorithms for the computation of PH. The two features that make PH appealing for applications are that it is computable via linear algebra and that it is stable with respect to perturbations in the measurement of data. In Section 5.5, we give a brief overview of stability results.




Figure 6: PH pipeline.

5.1 Data

As we mentioned in Section 1, types of data sets that one can study with PH include finite metric spaces, digital images, and networks. We now give a brief overview of how one can study these types of data sets using PH.

5.1.1 Networks

One can construe an undirected network as a 1-dimensional simplicial complex. If the network is weighted, then filtering by increasing or decreasing weight yields a filtered 1-dimensional simplicial complex. To obtain more refined information about the network, it is desirable to construct higher-dimensional simplices. There are various methods to do this. The simplest method, called a *weight rank clique filtration* (WRCF), consists of building a clique complex on each subnetwork. (See Section 5.2.1 for the definition of “clique complex.”) We refer the reader to [108] for an application of this method. Another method to study networks with PH consists of mapping the nodes of the network to points of a finite metric space. There are several ways to compute distances between nodes of a network; the method that we use in our benchmarking in Section 7 consists of computing a shortest path between nodes. For such a distance to be well-defined, note that one needs the network to be connected (although conventionally one takes the distance between nodes in different components to be infinity). There are many methods

to associate an unfiltered simplicial complex to both undirected and directed networks. See the book [72] for an overview of such methods, and see the paper [71] for an overview of PH for networks.

5.1.2 Digital images

As we mentioned in Section 1, digital images have a natural cubical structure: 2-dimensional digital images are made of pixels, and 3-dimensional images are made of voxels. Therefore, to study digital images, cubical complexes are more appropriate than simplicial complexes. Roughly, cubical complexes are spaces built from a union of vertices, edges, squares, cubes, and so on. One can compute PH for cubical complexes in a similar way as for simplicial complexes, and we will therefore not discuss this further in this paper. See [76] for a treatment of computational homology with cubical complexes rather than simplicial complexes and for a discussion of the relationship between simplicial and cubical homology. See [126] for an efficient algorithm and data structure for the computation of PH for cubical data, and [10] for an algorithm that computes PH for cubical data in an approximate way. For an application of PH and cubical complexes to movies, see [101].

Other approaches for studying digital images are also useful. In general, given a digital image that consists of N pixels or voxels, one can consider this image as a point in a $c \times N$ -dimensional space, with each coordinate storing a vector of length c representing the color of a pixel or voxel. Defining an appropriate distance function on such a space allows one to consider a collection of images (each of which has N pixels or voxels) as a finite metric space. A version of this approach was used in [24], in which the local structure of natural images was studied by selecting 3×3 patches of pixels of the images.

5.1.3 Finite metric spaces

As we mentioned in the previous two subsections, both undirected networks and image data can be construed as finite metric spaces. Therefore, methods to study finite metric spaces with PH apply to the study of networks and image data sets.

In some applications, points of a metric space have associated “weights.” For instance, in the study of molecules, one can represent a molecule as a union of balls in Euclidean space [131, 132]. For such data sets, one would therefore also consider a minimum filtration value (see Section 5.2 for the description of such filtration values) at which the point enters the filtration. In Table 2(g), we indicate which software libraries implement this feature.

5.2 Filtered Simplicial Complexes

In Section 3.2, we introduced the Čech complex, a classical simplicial complex from algebraic topology. However, there are many other simplicial complexes that are better suited for studying data from applications. We discuss them in this section.

To be a useful tool for the study of data, a simplicial complex has to satisfy some theoretical properties dictated by topological inference; roughly, if we build the simplicial complex on a set of points sampled from a space, then the homology of the simplicial complex has to approximate the homology of the space. For the Čech complex, these properties are guaranteed by the Nerve Theorem. Some of the complexes that we discuss in this subsection are motivated by a “sparsification paradigm”: they approximate the PH of known simplicial complexes but have fewer simplices than them. Others, like the Vietoris–Rips complex, are appealing because they can be computed efficiently. In this subsection, we also review reduction techniques, which are

heuristics that reduce the size of complexes without changing the PH. In Table 1, we summarize the simplicial complexes that we discuss in this subsection.

For the rest of this subsection (X, d) denotes a metric space, and S is a subset of X , which becomes a metric space with the induced metric. In applications, S is the collection of measurements together with a notion of distance, and we assume that S lies in the (unknown) metric space X . Our goal is then to compute persistent homology for a sequence of nested spaces $S_{\epsilon_1}, S_{\epsilon_2}, \dots, S_{\epsilon_l}$, where each space gives a “thickening” of S in X .

5.2.1 Vietoris–Rips complex

We have seen that one of the disadvantages of the Čech complex is that one has to check for a large number of intersections. To circumvent this issue, one can instead consider the Vietoris–Rips (VR) complex, which approximates the Čech complex. For a non-negative real number ϵ , the *Vietoris–Rips complex* $\text{VR}_\epsilon(S)$ at scale ϵ is defined as

$$\text{VR}_\epsilon(S) = \{\sigma \subseteq S \mid d(x, y) \leq 2\epsilon \text{ for all } x, y \in \sigma\}.$$

The sense in which the VR complex approximates the Čech complex is that, when S is a subset of Euclidean space, we have $\check{C}_\epsilon(S) \subseteq \text{VR}_\epsilon(S) \subseteq \check{C}_{\sqrt{2}\epsilon}(S)$. Deciding whether a subset $\sigma \subseteq S$ is in $\text{VR}_\epsilon(S)$ is equivalent to deciding if the maximal pairwise distance between any two vertices in σ is at most 2ϵ . Therefore, one can construct the VR complex in two steps. One first computes the ϵ -neighborhood graph of S . This is the graph whose vertices are all points in S and whose edges are

$$\{(i, j) \in S \times S \mid i \neq j \text{ and } d(i, j) \leq 2\epsilon\}.$$

Second, one obtains the VR complex by computing the *clique complex* of the ϵ -neighborhood graph. The clique complex of a graph is a simplicial complex that is defined as follows: The subset $\{x_0, \dots, x_k\}$ is a k -simplex if and only if every pair of vertices in $\{x_0, \dots, x_k\}$ is connected by an edge. Such a collection of vertices is called a *clique*. This construction makes it very easy to compute the VR complex, because to construct the clique complex one has only to check for pairwise distances — for this reason clique complexes are also called “lazy” in the literature. Unfortunately, the VR complex has the same worst-case complexity as the Čech complex. In the worst case, it can have up to $2^{|S|} - 1$ simplices and dimension $|S| - 1$.

In applications, one therefore usually only computes the VR complex up to some dimension $k \ll |S| - 1$. In our benchmarking, we often choose $k = 2$ and $k = 3$.

The paper [134] overviews different algorithms to perform both of the steps for the construction of the VR complex, and it introduces fast algorithms to construct the clique complex. For more details on the VR complex, see [51, 125]. For a proof of the approximation of the Čech complex by the VR complex, see [51]; see [79] for a generalization of this result.

5.2.2 The Delaunay complex

To avoid the computational problems of the Čech and VR complexes, we need a way to limit the number of simplices in high dimensions. The Delaunay complex gives a geometric tool to accomplish this task, and most of the new simplicial complexes that have been introduced for the study of data are based on variations of the Delaunay complex. The Delaunay complex and its dual, the Voronoi diagram, are central objects of study in computational geometry because they have many useful properties.

For the Delaunay complex, one usually considers $X = \mathbb{R}^d$, so we also make this assumption. We subdivide the space \mathbb{R}^d into regions of points that are closest to any of the points in S . More

precisely, for any $s \in S$, we define

$$V_s = \{x \in \mathbb{R}^d \mid d(x, s) \leq d(x, s') \text{ for all } s' \in S\}.$$

The collection of sets V_s is a cover for \mathbb{R}^d that is called the *Voronoi decomposition of \mathbb{R}^d with respect to S* , and the nerve of this cover is called the *Delaunay complex* of S and is denoted by $\text{Del}(S; \mathbb{R}^d)$. In general, the Delaunay complex does not have a geometric realization in \mathbb{R}^d . However, if the points S are “in general position”⁶ then the Delaunay complex has a geometric realization in \mathbb{R}^d that gives a triangulation of the convex hull of S . In this case, the Delaunay complex is also called the *Delaunay triangulation*.

The complexity of the Delaunay complex depends on the dimension d of the space. For $d \leq 2$, the best algorithms have complexity $\mathcal{O}(N \log N)$, where N is the cardinality of S . For $d \geq 3$, they have complexity $\mathcal{O}(N^{\lceil d/2 \rceil})$. The construction of the Delaunay complex is therefore costly in high dimensions, although there are efficient algorithms for the computation of the Delaunay complex for $d = 2$ and $d = 3$. Developing efficient algorithms for the construction of the Delaunay complex in higher dimensions is a subject of ongoing research. See [14] for a discussion of progress in this direction, and see [65] for more details on the Delaunay complex and the Voronoi diagram.

5.2.3 Alpha complex

We continue to assume that S is a finite set of points in \mathbb{R}^d . Using the Voronoi decomposition, one can define a simplicial complex that is similar to the Čech complex, but which has the desired property that (if the points S are in general position) its dimension is at most that of the space. Let $\epsilon > 0$, and let S_ϵ denote the union $\bigcup_{s \in S} B(s, \epsilon)$. For every $s \in S$, consider the intersection $V_s \cap B(s, \epsilon)$. The collection of these sets forms a cover of S_ϵ , and the nerve complex of this cover is called the *alpha (α) complex of S at scale ϵ* and is denoted by $A_\epsilon(S)$. The Nerve Theorem applies, and it therefore follows that $A_\epsilon(S)$ has the same homology as S_ϵ .

Furthermore, $A_\infty(S)$ is the Delaunay complex; and for $\epsilon < \infty$, the alpha complex is a subcomplex of the Delaunay complex. The alpha complex was introduced for points in the plane in [52], in 3-dimensional Euclidean space in [55], and for Euclidean spaces of arbitrary dimension in [48]. For points in the plane, there is a well-known speed-up for the alpha complex that uses a duality between 0-dimensional and 1-dimensional persistence for alpha complexes [50]. (See [84] for the algorithm, and see [83] for an implementation.)

5.2.4 Witness complexes

Witness complexes are very useful for analyzing large data sets, because they make it possible to construct a simplicial complex on a significantly smaller subset $L \subseteq S$ of points that are called “landmark” points. Meanwhile, because one uses information about all points in S to construct the simplicial complex, the points in S are called “witnesses.” Witness complexes can be construed as a “weak version” of Delaunay complexes. (See the characterization of the Delaunay complex in [37].)

Definition 6 *Let (S, d) be a metric space, and let $L \subseteq S$ be a finite subset. Suppose that σ is a non-empty subset of L . We then say that $s \in S$ is a weak witness for σ with respect to L if*

⁶A set S of points in \mathbb{R}^d is *in general position* if no $d + 2$ points of S lie on a d -dimensional sphere, and for any $d' < d$, no $d' + 2$ points of S lie on a d' -dimensional subspace that is isometric to $\mathbb{R}^{d'}$. In particular, a set of points S in \mathbb{R}^2 is in general position if no four points lie on a 2-dimensional sphere and no three points lie on a line.

and only if $d(s, a) \leq d(s, b)$ for all $a \in \sigma$ and for all $b \in L \setminus \sigma$. The weak Delaunay complex $\text{Del}^w(L; S)$ of S with respect to L has vertex set given by the points in L , and a subset σ of L is in $\text{Del}^w(L; S)$ if and only if it has a weak witness in S .

To obtain nested complexes, one can extend the definition of witnesses to ϵ -witnesses.

Definition 7 A point $s \in S$ is a weak ϵ -witness for σ with respect to L if and only if $d(s, a) \leq d(s, b) + \epsilon$ for all $a \in \sigma$ and for all $b \in L \setminus \sigma$.

Now we can define the *weak Delaunay complex* $\text{Del}^w(L; S, \epsilon)$ at scale ϵ to be the simplicial complex with vertex set L , and such that a subset $\sigma \subseteq L$ is in $\text{Del}^w(L; S, \epsilon)$ if and only if it has a weak ϵ -witness in S . By considering different values for the parameter ϵ , we thereby obtain nested simplicial complexes. The weak Delaunay complex is also called the “weak witness complex” or just the “witness complex” in the literature.

There is a modification of the witness complex called the *lazy witness complex* $\text{Del}_{\text{lazy}}^w(L; X, \epsilon)$. It is a clique complex, and it can therefore be computed more efficiently than the witness complex. The lazy witness complex has the same 1-skeleton as $\text{Del}^w(L; X, \epsilon)$, and one adds a simplex σ to $\text{Del}_{\text{lazy}}^w(L; X, \epsilon)$ whenever its edges are in $\text{Del}_{\text{lazy}}^w(L; X, \epsilon)$. Another type of modification of the witness complex yields *parametrized witness complexes*. Let $\nu = 1, 2, \dots$ and for all $s \in S$ define $m_\nu(s)$ to be the distance to the ν th closest landmark point. Furthermore, define $m_0(s) = 0$ for all $s \in S$. Let $W_\nu(L; S, \epsilon)$ be the simplicial complex whose vertex set is L and such that a 1-simplex $\sigma = \{x_0, x_1\}$ is in $W_\nu(L; X, \epsilon)$ if and only if there exists s in S for which

$$\max\{d(x_0, s), d(x_1, s)\} \leq m_\nu(s) + \epsilon.$$

A simplex σ is in $W_\nu(L; X, \epsilon)$ if and only if all of its edges belong to $W_\nu(L; X, \epsilon)$. For $\nu = 2$, note that $W_2(L; X, \epsilon) = \text{Del}_{\text{lazy}}^w(L; X, \epsilon)$. For $\nu = 0$, we have that $W_0(L; X, \epsilon)$ approximates the VR complex $\text{VR}(L; \epsilon)$. That is,

$$W_0(L; X, \epsilon) \subseteq \text{VR}(L; 2\epsilon) \subseteq W_0(L; X, 2\epsilon).$$

Note that parametrized witness complexes are often called “lazy witness complexes” in the literature, because they are clique complexes.

The weak Delaunay complex was introduced in [37], and parametrized witness complexes were introduced in [38]. Witness complexes can be rather useful for applications. Because their complexity depends on the number of landmark points, one can reduce complexity by computing simplicial complexes using a smaller number of vertices. However, there are theoretical guarantees for the witness complex only when S is the metric space associated to a low-dimensional Euclidean submanifold. It has been shown that witness complexes can be used to recover the topology of curves and surfaces in Euclidean space [4, 67], but they can fail to recover topology for submanifolds of Euclidean space of three or more dimensions [15]. Consequently, there have been studies of simplicial complexes that are similar to the witness complexes but with better theoretical guarantees. One of them is the “graph-induced complex” [44] (see Section 5.2.5).

5.2.5 Additional complexes

Many more complexes have been introduced for the fast computation of PH for large data sets. These include the graph-induced complex [44], which is a simplicial complex constructed on a subsample Q , and has better theoretical guarantees than the witness complex (see [74] for the companion software); an approximation of the VR complex that has a worst-case size that is linear in the number of data points [113]; an approximation of the Čech complex [79] whose worst-case size also scales linearly in the data; and an approximation of the VR complex via simplicial

Table 1: We summarize several types of complexes that are used for PH. We indicate the theoretical guarantees and the worst-case sizes of the complexes as functions of the cardinality N of the vertex set. For the witness complexes (see Section 5.2.4), L denotes the set of landmark points, while Q denotes the subsample set for the graph-induced complex (see Section 5.2.5).

Complex K	Size of K	Theoretical guarantee
Čech	$2^{\mathcal{O}(N)}$	Nerve theorem
Vietoris–Rips (VR)	$2^{\mathcal{O}(N)}$	Approximates Čech complex
Alpha (α)	$N^{\mathcal{O}(\lceil d/2 \rceil)}$ (N points in \mathbb{R}^d)	Nerve theorem
Witness	$2^{\mathcal{O}(L)}$	For curves and surfaces in Euclidean space
Graph-induced complex	$2^{\mathcal{O}(Q)}$	Approximates VR complex
Sparsified Čech	$\mathcal{O}(N)$	Approximates Čech complex
Sparsified VR	$\mathcal{O}(N)$	Approximates VR complex

collapses [46]. We do not discuss such complexes in detail, because thus far (at the time of writing) none of them have been implemented in publicly-available libraries for the computation of PH. (See Table 2 in Section 7 for information about which complexes have been implemented.)

5.2.6 Reduction techniques

Thus far, we have discussed techniques to build simplicial complexes with possibly “few” simplices. One can also take an alternative approach to speed up the computation of PH. For example, one can use a heuristic (i.e., a method without theoretical guarantees on the speed-up) to reduce the size of a filtered complex while leaving the PH unchanged.

For simplicial complexes, one such method is based on discrete Morse theory [111], which was adapted to filtrations of simplicial complexes in [96]. The basic idea of the algorithm developed in [96] is that one can compute a partial matching of the simplices in a filtered simplicial complex so that (i) pairs occur only between simplices that enter the filtration at the same step, (ii) unpaired simplices determine the homology, and (iii) one can remove paired simplices from the filtered complex without altering the total PH. Such deletions are examples of the elementary simplicial collapses that we mentioned in Section 3.1. Unfortunately, the problem of finding an optimal partial matching was shown to be NP complete [73], and one thus relies on heuristics to find partial matchings to reduce the size of the complex.

A method for the reduction of the size of a complex for clique complexes, such as the VR complex, was introduced in [135] and is called the *tidy-set method*. Using maximal cliques, this method extracts a minimal representation of the graph that determines the clique complex. Although the tidy-set method cannot be extended to filtered complexes, it can be used for the computation of zigzag PH (see Section 6) [136]. The tidy-set method is a heuristic, because it does not give a guarantee to minimize the size of the output complex.

5.3 From a Filtered Simplicial Complex to Barcodes

To compute the PH of a filtered simplicial complex K and obtain a barcode like the one illustrated in Fig. 5(c), we need to associate to it a matrix — the so-called *boundary matrix* — that stores information about the faces of every simplex. To do this, we place a total ordering on the simplices of the complex that is compatible with the filtration in the following sense:

- a face of a simplex precedes the simplex;

- a simplex in the i th complex K_i precedes simplices in K_j for $j > i$, which are not in K_i .

Let n denote the total number of simplices in the complex, and let $\sigma_1, \dots, \sigma_n$ denote the simplices with respect to this ordering. We construct a square matrix δ of dimension $n \times n$ by storing a 1 in $\delta(i, j)$ if the simplex σ_i is a face of simplex σ_j of codimension 1; otherwise, we store a 0 in $\delta(i, j)$.

Once one has constructed the boundary matrix, one has to reduce it using Gaussian elimination.⁷ In the following subsections, we discuss several algorithms for reducing the boundary matrix.

5.3.1 Standard algorithm

The so-called **standard algorithm** for the computation of PH was introduced for the field \mathbb{F}_2 in [53] and for general fields in [137]. For every $j \in \{1, \dots, n\}$, we define $\text{low}(j)$ to be the largest index value i such that $\delta(i, j)$ is different from 0.⁸ If column j only contains 0 entries, then the value of $\text{low}(j)$ is undefined. We say that the boundary matrix is *reduced* if the map low is injective on its domain of definition. In Alg. 1, we illustrate the standard algorithm for reducing the boundary matrix. Because this algorithm operates on columns of the matrix from left to right, it is also sometimes called the “column algorithm.” In the worst case, the complexity of the standard algorithm is cubic in the number of simplices.

Algorithm 1 The standard algorithm for the reduction of the boundary matrix to barcodes.

```

for  $j = 1$  to  $n$  do
  while there exists  $i < j$  with  $\text{low}(i) = \text{low}(j)$  do
    add column  $i$  to column  $j$ 
  end while
end for

```

5.3.2 Reading off the intervals

Once the boundary matrix is reduced, one can read off the intervals of the barcode by pairing the simplices in the following way:

- If $\text{low}(j) = i$, then the simplex σ_j is paired with σ_i , and the entrance of σ_i in the filtration causes the birth of a feature that dies with the entrance of σ_j .
- If $\text{low}(j)$ is undefined, then the entrance of the simplex σ_j in the filtration causes the birth of a feature. If there exists k such that $\text{low}(k) = j$, then σ_j is paired with the simplex σ_k , whose entrance in the filtration causes the death of the feature. If no such k exists, then σ_j is unpaired.

A pair (σ_i, σ_j) gives the half-open interval $[\text{dg}(\sigma_i), \text{dg}(\sigma_j))$ in the barcode, where for a simplex $\sigma \in K$ we define $\text{dg}(\sigma)$ to be the smallest number l such that $\sigma \in K_l$. An unpaired simplex σ_k gives the infinite interval $[\text{dg}(\sigma_k), \infty)$. We give an example of PH computation in Fig. 7.


⁷As we mentioned in Section 4, for the reduction of the boundary matrix and thus the computation of PH, it is crucial that one uses simplicial homology with coefficients in a field; see [137] for details.

⁸This map is called “low” in the literature, because one can think of it as indicating the index of the “lowest” row — the one that is nearest to the bottom of the page on which one writes the boundary matrix — that contains a 1 in column j .

(a) A filtered simplicial complex:



(b) We put a total order on the simplices that is compatible with the filtration:



where σ_i denotes the i th simplex in this order.

(c) (Left) The boundary matrix B for the filtered simplicial complex in (a) with respect to order on simplices in (b), and (right) its reduction \bar{B} given by applying Alg. 1 (one first adds column 5 to column 6, and then column 4 to column 6):

$$B = \begin{pmatrix} 0 & 0 & 0 & 1 & 1 & 0 & 0 \\ 0 & 0 & 0 & 1 & 0 & 1 & 0 \\ 0 & 0 & 0 & 0 & 1 & 1 & 0 \\ 0 & 0 & 0 & 0 & 0 & 0 & 1 \\ 0 & 0 & 0 & 0 & 0 & 0 & 1 \\ 0 & 0 & 0 & 0 & 0 & 0 & 1 \\ 0 & 0 & 0 & 0 & 0 & 0 & 0 \end{pmatrix} \quad \bar{B} = \begin{pmatrix} 0 & 0 & 0 & 1 & 1 & 0 & 0 \\ 0 & 0 & 0 & 1 & 0 & 0 & 0 \\ 0 & 0 & 0 & 0 & 1 & 0 & 0 \\ 0 & 0 & 0 & 0 & 0 & 0 & 1 \\ 0 & 0 & 0 & 0 & 0 & 0 & 1 \\ 0 & 0 & 0 & 0 & 0 & 0 & 1 \\ 0 & 0 & 0 & 0 & 0 & 0 & 0 \end{pmatrix}.$$

(d) We read off the following intervals from the matrix \bar{B} in (c):

- σ_1 is positive, unpaired; this gives the interval $[1, \infty)$ in H_0 .
- σ_2 is positive, paired with σ_4 ; this gives no interval, because σ_2 and σ_4 enter at the same time in the filtration.
- σ_3 is positive, paired with σ_5 ; this gives the interval $[2, 3)$ in H_0 .
- σ_6 is positive, paired with σ_7 ; this gives the interval $[3, 4)$ in H_1 .

Figure 7: Example of PH computation with the **standard algorithm** (see Alg. 1).

5.3.3 Other algorithms

After the introduction of the standard algorithm, several new algorithms were developed. Each of these algorithms gives the same output for the computation of PH, so we only give a brief overview and references to these algorithms, as one does not need to know them to compute PH with one of the publicly-available software packages. In Section 7.2, we indicate which implementation of these libraries is best suited to which data set.

As we mentioned in Section 5.3.1, in the worst case, the standard algorithm has cubic complexity in the number of simplices. This bound is sharp, as Morozov gave an example of a complex with cubic complexity in [98]. Note that in cases such as when matrices are sparse, complexity is less than cubic. Milosavljević, Morozov, and Skraba [95] introduced an algorithm for the reduction of the boundary matrix in $\mathcal{O}(n^\omega)$, where ω is the matrix-multiplication coefficient (i.e., $\mathcal{O}(n^\omega)$ is the complexity of the multiplication of two square matrices of size n). At present, the best bound for ω is 2.376 [33]. Many other algorithms have been proposed for the reduction of the boundary matrix. These algorithms give a heuristic speed-up for many data sets and complexes (see the benchmarkings in the forthcoming references), but they still have cubic complexity in the number of simplices. Sequential algorithms include the **twist algorithm** [28] and the **dual algorithm** [40,41]. (Note that the dual algorithm is known to give a speed-up when one computes PH with the VR complex, but not necessarily for other types of complexes (see also the results of our benchmarking for the **vertebra** data set in the SI).) Parallel algorithms in a shared setting include the **spectral-sequence algorithm** [51, Section VII.4] and the **chunk algorithm** [7], and parallel algorithms in a distributed setting include the **distributed algorithm** [8]. The **multifield algorithm** is a sequential algorithm that allows the simultaneous computation of PH over several fields [16].

5.4 Statistical interpretation of topological summaries

Once we have obtained the barcodes, we need to interpret the results of our computations. In applications, one often wants to compare the output of a computation for a certain data set with the output for a null model. Alternatively, one may be studying data sets from the output of a generative model (e.g., many realizations from a model of random networks), and it is then necessary to average results over multiple realizations. In the first instance, one needs both a way to compare the two different outputs and a way to evaluate the significance of the result for the original data set. In the second case, one needs a way to calculate appropriate averages (e.g., summary statistics) of the result of the computations.

If one wants to use PH in applications, one thus needs a reliable way to apply statistical methods to the output of the computation of PH. To our knowledge, statistical methods for PH were addressed for the first time in the paper [20]. Roughly, there are two basic types of current approaches to this problem. In the first approach, one studies properties of a metric space whose points are persistence diagrams, while in the second approach one studies “features” of persistence diagrams.

More precisely, in the first approach one considers an appropriately defined “space of persistence diagrams,” defines a distance function on it, studies geometric properties of this space, and does standard statistical calculations (means, medians, statistical tests, and so on). Recall that a persistence diagram (see Fig. 5 for an example) is a multiset of points in \mathbb{R}^2 and that it gives the same information as a barcode. We now give the following precise definition of a persistence diagram.

Definition 8 *A persistence diagram is a multiset that is the union of a finite multiset of points in \mathbb{R}^2 with the multiset of points on the diagonal $\Delta = \{(x, y) \in \mathbb{R}^2 \mid x = y\}$, where each point on*

the diagonal has infinite multiplicity.

In this definition, we include all of the points on the diagonal in \mathbb{R}^2 with infinite multiplicity. We include the diagonal points for technical reasons; roughly, it is desirable to be able to compare persistence diagrams by studying bijections between their points, and persistence diagrams must thus be sets with the same cardinality.

Given two persistence diagrams X and Y , we consider the following general definition of distance between X and Y .

Definition 9 Let $p \in [1, \infty]$. The p th Wasserstein distance between X and Y is defined as

$$W_p[d](X, Y) := \inf_{\phi: X \rightarrow Y} \left[\sum_{x \in X} d[x, \phi(x)]^p \right]^{1/p}$$

for $1 \leq p < \infty$ and as

$$W_\infty[d](X, Y) := \inf_{\phi: X \rightarrow Y} \sup_{x \in X} d[x, \phi(x)]$$

for $p = \infty$, where d is a metric on \mathbb{R}^2 and ϕ ranges over all bijections from X to Y .

Usually, one takes $d = L_p$ for $1 \leq p \leq \infty$. One of the most commonly employed distance functions is the *bottleneck distance* $W_\infty[L_\infty]$.

The development of statistical analysis on the space of persistence diagrams is an area of ongoing research, and at present there are few tools that can be used in applications. We refer the reader to the publications [94, 99, 122] for research in this direction.

The second approach for the development of statistical tools for PH consists of mapping the space of persistence diagrams to spaces (e.g., Banach spaces) that are amenable to statistical analysis. Such methods include persistence landscapes [17] (see also the tutorial [19]), the space of algebraic functions [3], and persistence images [29]. See the papers [80, 117] for applications of persistence landscapes.

The library DIONYSUS [97] implements the computation of the bottleneck and Wasserstein distance (for $d = L_\infty$), while the PERSISTENCE LANDSCAPE TOOLBOX [47] and the TDA PACKAGE [58] implement many tools for the statistical interpretation of persistent homology.

5.5 Stability

As we mentioned in Section 1, PH is useful for applications because it is stable with respect to small perturbations in the input data.

The first stability theorem for PH, proven in [32], asserts that, under favorable conditions, step (2) in the pipeline in Fig. 6 is 1-Lipschitz with respect to suitable distance functions on filtered complexes and the bottleneck distance for barcodes (see Section 5.4). This result was generalized in the papers [18, 21, 27]. Stability for PH is an active area of research; for an overview of stability results, their history and recent developments, see [104, Chapter 3].

6 Excursus: Generalized persistence

One can use the algorithms that we described in Section 5 to compute PH when one has a sequence of complexes with inclusion maps that are all going in the same direction, as in the following diagram:

$$\cdots \rightarrow K_{i-1} \rightarrow K_i \rightarrow K_{i+1} \rightarrow \cdots$$

An algorithm, called the **zigzag algorithm**, for the computation of PH for inclusion maps that do not all go in the same direction, as, e.g., in the diagram

$$\cdots \rightarrow K_{i-1} \rightarrow K_i \leftarrow K_{i+1} \rightarrow \cdots$$

was introduced in [23]. In the more general setting in which maps are not inclusions, one can still compute PH using the **simplicial map algorithm** [45].

One may also wish to vary two or more parameters instead of one. This yields multi-filtered simplicial complexes, as, e.g., in the following diagram:

$$\begin{array}{ccccccc} & & \vdots & & \vdots & & \vdots \\ & & \uparrow & & \uparrow & & \uparrow \\ \cdots & \rightarrow & K_{j+1,i-1} & \rightarrow & K_{j+1,i} & \rightarrow & K_{j+1,i+1} \rightarrow \cdots \\ & & \uparrow & & \uparrow & & \uparrow \\ \cdots & \rightarrow & K_{j,i-1} & \rightarrow & K_{j,i} & \rightarrow & K_{j,i+1} \rightarrow \cdots \\ & & \uparrow & & \uparrow & & \uparrow \\ \cdots & \rightarrow & K_{j-1,i-1} & \rightarrow & K_{j-1,i} & \rightarrow & K_{j-1,i+1} \rightarrow \cdots \\ & & \uparrow & & \uparrow & & \uparrow \\ & & \vdots & & \vdots & & \vdots \end{array}$$

In this case, one speaks of *multi-parameter persistent homology*. Unfortunately, the Fundamental Theorem of Persistent Homology is no longer valid if one filters with more than one parameter, and there is no such thing as a “generalized interval.” The topic of multi-parameter persistence is under active research, and several approaches are being studied to extract topological information from multi-filtered simplicial complexes. See [25, 104] for the theory of multi-parameter persistent homology, and see [89] (and [88] for its companion paper) for upcoming software for the visualization of 2-parameter persistent homology.

7 Software

There are several publicly-available implementations for the computation of PH. We give an overview of the libraries with accompanying peer-reviewed publication and summarize their properties in Table 2.

The software package JAVAPLEX [118], which was developed by the computational topology group at Stanford University, is based on the PLEX library [107], which to our knowledge is the first piece of software to implement the computation of PH. PERSEUS [100] was developed to implement Morse-theoretic reductions [96] (see Section 5.2.6). JHOLES [12] is a Java library for computing the weight rank clique filtration for weighted undirected networks [108]. DIONYSUS [97] is the first software package to implement the **dual algorithm** [40, 41]. PHAT [9] is a library that implements several algorithms and data structures for the fast computation of barcodes, takes a boundary matrix as input, and is the first software to implement a matrix-reduction algorithm that can be executed in parallel. DIPHA [6], a spin-off of PHAT, implements a distributed computation of the matrix-reduction algorithm. GUDHI [93] implements new data structures for simplicial complexes and the boundary matrix. It also implements the **multi-field algorithm**, which allows simultaneous computation of PH over several fields [16]. This library is currently under intense development. The library RIPSER [5], the most recently developed software of the set that we examine, uses several optimizations and shortcuts to speed up the computation of PH with the VR complex. This library does not have an accompanying peer-reviewed publication.

However, because it is currently the best-performing (both in terms of memory usage and in terms of wall-time seconds⁹) library for the computation of PH with the VR complex, we include it in our study. The library SIMPPERS [75] implements the **simplicial map algorithm**. Libraries that implement techniques for the statistical interpretation of barcodes include the TDA PACKAGE [58] and the PERSISTENCE LANDSCAPE TOOLBOX [47]. RIVET, a package for visualizing 2-parameter persistent homology, is slated to be released soon [89]. We summarize the properties of the libraries that we mentioned in this paragraph in Table 2, and we discuss the performance for a selection of them in Section 7.1.3 and in the SI. For a list of programs, see <https://github.com/n-otter/PH-roadmap>.

7.1 Benchmarking

We benchmark a subset of the currently available open-source libraries with peer-reviewed publication for the computation of PH. To our knowledge, the published open-source libraries are JHOLES, JAVAPLEX, PERSEUS, DIONYSUS, PHAT, DIPHA, SIMPPERS, and GUDHI. To these, we add the library RIPSER, which is currently the best-performing library to compute PH with the VR complex. To study the performance of the packages, we restrict our attention to the algorithms that are implemented by the largest number of libraries. These are the VR complex and the standard and dual algorithms for the reduction of the boundary matrix. PHAT only takes a boundary matrix as input, so it is not possible to conduct a direct comparison of it with the other implementations. However, the fast data structures and algorithms implemented in PHAT are also implemented in its spin-off software DIPHA, which we include in the benchmarking. The software JHOLES computes PH using the WRCF for weighted undirected networks, and SIMPPERS takes a map of simplicial complexes as input, so these two libraries cannot be compared directly to the other libraries. In the SI, we report benchmarking of some additional features that are implemented by some of the six libraries (i.e., JAVAPLEX, PERSEUS, DIONYSUS, DIPHA, GUDHI, and RIPSER) that we test. Specifically, we report results for the computation of PH with cubical complexes for image data sets and the computation of PH with witness, alpha, and Čech complexes.

We study the software packages JAVAPLEX, PERSEUS, DIONYSUS, DIPHA, GUDHI, and RIPSER using both synthetic and real-world data from three different perspectives:

1. Performance measured in CPU seconds and wall-time (i.e., elapsed time) seconds.
2. Memory required by the process.
3. Maximum size of simplicial complex allowed by the software.

7.1.1 Data sets

In this subsection, we describe the data sets that we use for our benchmarking. We use data sets from a variety of different mathematical and scientific areas and applications. In each case, when possible, we use data sets that have already been studied using PH. Our list of data sets is far from complete; we view this list as an initial step towards building a comprehensive collection of benchmarking data sets for PH.

Data sets (1)–(4) are synthetic: these arise from topology (1), stochastic topology (2), dynamical systems (3), and from an area at the intersection of network theory and fractal geometry (4) and which was first used to study connection patterns of the cerebral cortex (see below for details). Data sets (5)–(12) are from empirical experiments and measurements: they arise from

⁹“Wall time” is the amount of elapsed time perceived by a human.

Table 2: Overview of existing software for the computation of PH that have an accompanying peer-reviewed publication (and also RIPSER [5], because of its performance). The symbol “–” signifies that the associated feature is not implemented. For each software package, we indicate the following items. (a) The language in which it is implemented. (b) The implemented algorithms for the computation of barcodes from the boundary matrix. (c) The coefficient fields for which PH is computed, where the letter p denotes any prime number in the coefficient field \mathbb{F}_p . (d) The type of homology computed. (e) The filtered complexes that are computed, where VR stands for Vietoris–Rips complex, W stands for the weak witness complex, W_ν stands for parametrized witness complexes, WRCF stands for the weight rank clique filtration, and α stands for the alpha complex. PERSEUS, DIPHA, and GUDHI implement the computation of the lower-star filtration [54] of a weighted cubical complex; one inputs data in the form of a d -dimensional array; the data is then interpreted as a d -dimensional cubical complex, and its lower-star filtration is computed. (See the Tutorial in the SI for more details.) Note that DIPHA and GUDHI use the efficient representation of cubical complexes presented in [126], so the size of the cubical complex that is computed by these libraries is smaller than the size of the resulting complex with PERSEUS. (f) The filtered complexes that one can give as input. JAVAPLEX supports the input of a filtered CW complex for the computation of cellular homology [69]; in contrast with simplicial complexes, there do not currently exist algorithms to assign a cell complex to point-cloud data. (g) Additional features implemented by the library. JAVAPLEX supports the computation of some constructions from homological algebra (see [118] for details), and PERSEUS implements the computation of PH with the VR for points with different “birth times” (see Section 5.1.3). The library DIONYSUS implements the computation of vineyards [23] and circle-valued functions [41]. Both JAVAPLEX and DIONYSUS support the output of representatives of homology classes for the intervals in a barcode. (i) Whether visualization of the output is provided.

Software	JAVAPLEX	PERSEUS	JHOLES	DIONYSUS	PHAT	DIPHA	GUDHI	SIMPERS	RIPSER
(a) Language	Java	C++	Java	C++	C++	C++	C++	C++	C++
(b) Algorithms for PH	standard, dual, zigzag	Morse reductions, standard	standard (uses JAVAPLEX)	standard, dual, zigzag	standard, dual, twist, chunk, spectral seq.	twist, dual, distributed	dual, multifield	simplicial map	twist, dual
(c) Coeff. field	\mathbb{Q} , \mathbb{F}_p	\mathbb{F}_2	\mathbb{F}_2	\mathbb{F}_2 (standard, zigzag), \mathbb{F}_p (dual)	\mathbb{F}_2	\mathbb{F}_2	\mathbb{F}_p	\mathbb{F}_2	\mathbb{F}_p
(d) Homology	simplicial, cellular	simplicial, cubical	simplicial	simplicial	simplicial, cubical	simplicial, cubical	simplicial, cubical	simplicial	simplicial
(e) Filtrations computed	VR, W, W_ν	VR, lower star of cubical complex	WRCF	VR, α , Čech	–	VR, lower star of cubical complex	VR, α , W lower star of cubical complex	–	VR
(f) Filtrations as input	simplicial complex, zigzag, CW	simplicial complex, cubical complex	–	simplicial complex, zigzag	boundary matrix of simpl. complex	boundary matrix of simpl. complex	–	map of simpl. complexes	–
(g) Additional features	Computes some hom. alg. constructions, homology generators	weighted points for VR	–	vineyards, circle-valued functions, homology generators	–	–	–	–	–
(h) Visualization	barcodes	persistence diagram	–	–	–	persistence diagram	–	–	–

phylogenetics (5)–(6), genomics (9), neuroscience (8), image analysis (7), medical imaging (10), political science (11), and scientometrics (12).

In each case, these data sets are of one of the following three types: point clouds, weighted undirected networks, and grey-scale digital images. To obtain a point cloud from a real-world weighted undirected network, we compute shortest paths using the inverse of the nonzero weights on edges as distances between nodes (except for the US Congress networks; see below). For the synthetic networks, the values assigned to edges are interpreted as distances between nodes, and we therefore use these values to compute shortest paths. We make all processed versions of the data sets that we use in the benchmarking available at https://github.com/n-otter/PH-roadmap/tree/master/data_sets. We provide the scripts that we used to produce the synthetic data sets at https://github.com/n-otter/PH-roadmap/tree/master/matlab/synthetic_data_sets_scripts.

We now describe all data sets in detail:

- (1) **Klein bottle.** The Klein bottle is a one-sided nonorientable surface (see Fig. 8). We linearly sample points from the Klein bottle using its “figure-8” immersion in \mathbb{R}^3 and size sample of 400 points. We denote this data set by **Klein**. Note that the image of the immersion of the Klein bottle does not have the same homotopy type as the original Klein bottle, but they do have the same singular homology¹⁰ with \mathbb{F}_2 coefficients. We have $H_0(B) = \mathbb{F}_2$, $H_1(B) = \mathbb{F}_2 \oplus \mathbb{F}_2$, and $H_2(B) = \mathbb{F}_2$, where B denotes the Klein bottle and $H_i(B)$ is the i th singular homology group with \mathbb{F}_2 coefficients.
- (2) **Random VR complexes (uniform distribution)** [77]. The parameters for this model are positive integers N and d ; the random VR complex for parameters N and d is the VR complex $\text{VR}_\epsilon(X)$, where X is a set of N points sampled from \mathbb{R}^d . (Equivalently, the random VR complex is the clique complex on the random geometric graph $G(N, \epsilon)$ [105].) We sample N points uniformly at random from $[0, 1]^d$. We choose $(N, d) = (50, 16)$ and we denote this data set by **random**. The homology of random VR complexes was studied in [77].
- (3) **Vicsek biological aggregation model.** This model was first introduced in [124] and was studied using PH in [121]. We implement the model in the form in which it appears in [121]. The model describes the motion of a collection of particles that interact in a square with periodic boundary conditions. The parameters for the model are the length l of the side of the square, the initial angle θ_0 , the fixed absolute value for the velocity v_0 , the number of particles N , a noise parameter η , and the number T of time steps. The output of the model is a point cloud in 3-dimensional Euclidean space in which each point is specified by its position in the 2-dimensional box and its velocity angle. We run three simulations of the model using the parameter values used in [121]. For each simulation, we choose two point clouds that correspond to two different time frames. See [121] for further details. We denote this data set by **Vicsek**.
- (4) **Fractal networks.** These are self-similar networks introduced in [116] to investigate whether connection patterns of the cerebral cortex are arranged in self-similar patterns. The parameters for this model are natural numbers b , k , and n . To generate a fractal network, one starts with a fully-connected network on 2^b nodes. Two copies of this network are connected to each other so that the “connection density” between them is k^{-1} , where the connection density is the number of edges between the two copies divided by the number of

¹⁰*Singular homology* is a method that assigns to every topological space homology groups encoding invariants of the space, in an analogous way as simplicial homology assigns homology groups to simplicial complexes. We refer the reader to [69] for an account of singular homology.

total possible edges between them. Two copies of the resulting network are connected with connection density k^{-2} . One repeats this type of connection process until the network has size 2^n , but with a decrease in the connection density by a factor of $1/k$ at each step.

We define distances between nodes in two different ways: (1) uniformly at random, and (2) with linear weight-degree correlations. In the latter, the distance between nodes i and j is distributed as $k_i k_j X$, where k_i is the degree of node i and X is a random variable uniformly distributed on the unit interval. We use the parameters $b = 5$, $n = 9$, and $k = 2$; and we compute PH for the weighted network and for the network in which all adjacent nodes have distance 1. We denote this data set by **fract** and distinguish between the two ways of defining distances between weights using the abbreviations “r” for random, and “l” for linear.

- (5) Genomic sequences of the HIV virus. We construct a finite metric space using the independent and concatenated sequences of the three largest genes — **gag**, **pol**, and **env** — of the HIV genome. We take 1088 different genomic sequences and compute distances between them by using the Hamming distance. We use the aligned sequences studied using PH in [26]. (The authors of that paper retrieved the sequences from [85].) We denote this data set by **HIV**.
- (6) Genomic sequences of H3N2. These are 1000 different genomic sequences of H3N2 influenza. We compute the Hamming distance between sequences. We use the aligned sequences studied using PH in [26]. We denote this data set by **H3N2**.
- (7) Stanford Dragon graphic. We sample points uniformly at random from 3-dimensional scans of the dragon [86], whose reconstruction we show in Fig. 8. The sample sizes contain 1000 and 2000 points. We denote these data sets by **drag 1** and **drag 2**, respectively.
- (8) *C. elegans* neuronal network [108]. A weighted, undirected network in which each node is a neuron and edges represent synapses or gaps junctions. We denote this data set by **eleg**.
- (9) Human genome. A weighted, undirected network representing a sample of the human genome. We use the network studied using PH in [108]. (The authors of that paper created the sample using data retrieved from [36].) Each node represents a gene, and weighted edges between nodes represent the correlation of the expression level of two genes. We denote this data set by **genome**.
- (10) Grey-scale image: 3-dimensional rotational angiography scan of a head with an aneurysm. This data set was used in the benchmarking in [8]. This data set is given by a 3-dimensional array of size $512 \times 512 \times 512$, with each entry storing an integer that represents the grey value for the corresponding voxel. We retrieved the data set from the repository [1]. We denote this data set by **vertebra**.
- (11) US Congress roll-call voting networks. These two networks (the Senate and House of Representatives from the 104th United States Congress) are constructed using the procedure in [128] from data compiled by [110]. In each network, a node is a legislator (Senators in one data set and Representatives in the other), and there is a weighted edge between legislators i and j , where the weight $w_{i,j}$ is a number in $[0, 1]$ (it is equal to 0 if and only if legislators i and j never voted the same way on any bill) given by the number of times the two legislators voted in the same way divided by the total number of bills on which they both voted. See [128] for further details. We denote the networks from the Senate and House by **senate** and **house**, respectively. The network **senate** has 103 nodes, and the

network **house** has 445 nodes. To compute shortest paths, we define the distance between two nodes i and j to be $1 - w_{i,j}$. In the 104th Congress, no two politicians voted in the same way on every bill, so we do not have distinct nodes with 0 distance between them. (This is important, for example, if one wants to apply multidimensional scaling.)

- (12) Network of network scientists. This is a weighted undirected network representing the largest connected component of a collaboration network of network scientists [102]. Nodes represent authors and edges represent collaborations, where weights indicate the number of joint papers. The largest connected components consists of 379 nodes. We denote this data set by **netw-sc**.




Figure 8: Plot of the image of the figure-8 immersion of the Klein bottle and the reconstruction of the Stanford Dragon (retrieved from [86]).

7.1.2 Machines and compilers

We tested the libraries on both a cluster and a shared-memory system. The cluster is a Dell Sandybridge cluster, it has 1728 (i.e., 108×16) cores of 2.0GHz Xeon SandyBridge, RAM of 64 GiB in 80 nodes and RAM of 128 GiB in 4 nodes, and a scratch disk of 20 TB. It runs the operating system (OS) Red Hat Enterprise Linux 6. The shared-memory system is an IBM System x3550 M4 server with 16 (i.e., 2×8) cores of 3.3GHz, RAM of 768 GB, and storage of 3 TB. It runs the OS Ubuntu 14.04.01.¹¹ The major difference in running shared algorithms on the shared-memory system versus the distributed-memory system is that each node in the former has much more available RAM than in the latter. (See also the difference in performance between computations on cluster and shared memory system in Tables 3 and 4.) To compile GUDHI, DIPHA, PERSEUS and DIONYSUS we used the compiler `gcc` 4.8.2 on the cluster, and `gcc` 4.8.4 on the shared-memory system; for both machines we used the (default) optimization `-O3`. Additionally, we used `openmpi` 1.8.3 for DIPHA.

7.1.3 Tests and results

We now report the details and results of the tests that we performed. We have made the data sets, header file to measure memory, and other information related to the tests available at <https://github.com/n-otter/PH-roadmap>. Of the six software packages that we study, four implement the computation of the dual algorithm, and four implement the standard algorithm. It is reported in [118] that JAVAPLEX implements the dual algorithm, but the implementation of the algorithm has a bug and gives a wrong output. To our knowledge, this bug has not yet been fixed (at the time of writing), and we therefore test only the standard algorithm.

For the computations on the cluster, we compare the libraries running both the dual algorithm and the standard algorithm. The package DIPHA is the only one to implement a distributed

¹¹Note that we performed the computations for GUDHI and RIPSER at a different point in time, during which the shared-memory system was running the OS Ubuntu 16.04.01.

computation. As a default, we run the software on one node and 16 cores; we only increase the number of nodes and cores employed when the machine runs out of memory. However, augmenting the number of nodes can make the computations faster (in terms of CPU seconds) for complexes of all sizes.¹² We see this in our experiments, and it is also discussed in [8]. For the other packages, we run the computations on a single node with one core.

For computations on the shared-memory system, we compare the libraries using only the dual algorithm if they implement it, and we otherwise use the standard algorithm. For the shared-memory system, we run all packages (including DIPHA) on a single core.

In our benchmarking, we report mean computation times and memory measurements. In Table 3, we give the computation times for the different software packages. We measure elapsed and CPU time by using the `time` function in Linux. We report computation times with a precision of one second; if a computation took only fractions of a second, we report “one second” as the computation time. For space reasons, we report results for a subset of the computations, and we refer the reader to the SI for a tabulation of the rest of our computations. In Table 4, we report the memory used by the processes in terms of maximum resident set size (RSS); in other words, we give the maximum amount of real RAM a program has used during its execution. We measure the maximum RSS using the `getrusage` function in Linux. The header file that we use to measure memory is available at <https://github.com/n-otter/PH-roadmap>. In DIPHA, the measurement of memory is already implemented by the authors of the software. They also use the `getrusage` function in Linux. The package JAVAPLEX is written in Java, and we thus cannot measure its memory as we do for the other packages. However, one can infer memory requirements for this software package using the value of the maximal heap size necessary to perform the computations; we report this value in Table 4. In Table 5, we give the maximum size of the simplicial complex for which we were able to compute PH with each software package in our benchmarkings.

7.2 Conclusions from our benchmarking

Our tests suggest that RIPSER is the best-performing library currently available for the computation of PH with the Vietoris–Rips complex, and in order of decreasing performance, that GUDHI and DIPHA are the next-best implementations. For the computation of PH with cubical complexes, GUDHI outperforms DIPHA by a factor of 3 to 4 in terms of memory usage, and DIPHA outperforms GUDHI in terms of wall-time seconds by a factor of 1 to 2 (when running on one core on a shared-memory system). Both DIPHA and GUDHI significantly outperform the implementation in PERSEUS. For the computation of PH with the alpha complex, we did not observe any significant differences in performance between the libraries GUDHI and DIONYSUS. Because the alpha complex has fewer simplices than the other complexes that we used in our tests, further tests with larger data sets may be appropriate in future benchmarkings.

There is a huge disparity between implementations of the dual and standard algorithms. In our benchmarking, the dual implementations outperformed standard ones both in terms of computation time (with respect to both CPU and wall-time seconds) and in terms of the amount of memory used. This significant difference in performance and memory usage was also revealed for the software package DIONYSUS in [40].

To conclude, in our benchmarking, the fastest software packages were RIPSER, GUDHI, and DIPHA. For small complexes, the software packages PERSEUS and JAVAPLEX are good choices, because they are the easiest ones to use. (They are the only libraries that need only to be

¹²Based on the results of our tests, we think of small, medium, and large complexes, respectively, as complexes with a size of order of magnitude of up to 10 million simplices, between 10 million and 100 million simplices, and between 100 million and a billion simplices.

Table 3: Performance of the software packages measured in wall-time (i.e., elapsed time), and CPU seconds (for the computations running on the cluster). For each data set, we indicate the size of the simplicial complex and the maximum dimension up to which we construct the VR complex. For all data sets, we construct the filtered VR complex up to the maximum distance between any two points. We indicate the implementation of the standard algorithm using the abbreviation “st” following the name of the package, and we indicate the implementation of the dual algorithm using the abbreviation “d.” The symbol “-” signifies that we were unable to finish computations for this data set, because the machine ran out of memory. PERSEUS implements only the standard algorithm, and GUDHI and RIPSER implement only the dual algorithm. (a,b) We run DIPHA on one node and 16 cores for the data sets **eleg**, **Klein**, and **genome**; on 2 nodes of 16 cores for the **HIV** data set; on 2 and 3 nodes of 16 cores for the dual and standard implementations, respectively, for **drag 2**; and on 8 nodes of 16 cores for **random**. (The maximum number of processes that we could use at any one time was 128.) (c) We run DIPHA on a single core.

(a) Computations on cluster: wall-time seconds

Data set	eleg	Klein	HIV	drag 2	random	genome
Size of complex	4.4×10^6	1.1×10^7	2.1×10^8	1.3×10^9	3.1×10^9	4.5×10^8
Max. dim.	2	2	2	2	8	2
JAVAPLEX (st)	84	747	-	-	-	-
DIONYSUS (st)	474	1830	-	-	-	-
DIPHA (st)	6	90	1631	142559	-	9110
PERSEUS	543	1978	-	-	-	-
DIONYSUS (d)	513	145	-	-	-	-
DIPHA (d)	4	6	81	2358	5096	232
GUDHI	36	89	1798	14368	-	4753
RIPSER	1	1	2	6	349	3

(b) Computations on cluster: CPU seconds

Data set	eleg	Klein	HIV	drag 2	random	genome
Size of complex	4.4×10^6	1.1×10^7	2.1×10^8	1.3×10^9	3.1×10^9	4.5×10^8
Max. dim.	2	2	2	2	8	2
JAVAPLEX (st)	284	1031	-	-	-	-
DIONYSUS (st)	473	1824	-	-	-	-
DIPHA (st)	68	1360	25950	1489615	-	130972
PERSEUS	542	1974	-	-	-	-
DIONYSUS (d)	513	145	-	-	-	-
DIPHA (d)	39	73	1276	37572	79691	3622
GUDHI	36	88	1794	14351	-	4764
RIPSER	1	1	2	5	348	2

(c) Computations on shared-memory system: wall-time seconds

Data set	eleg	Klein	HIV	drag 2	genome	fract r
Size of complex	3.2×10^8	1.1×10^7	2.1×10^8	1.3×10^9	4.5×10^8	2.8×10^9
Max. dim.	3	2	2	2	2	3
JAVAPLEX (st)	13607	1358	43861	-	28064	-
PERSEUS	-	1271	-	-	-	-
DIONYSUS (d)	-	100	142055	35366	-	572764
DIPHA (d)	926	13	773	4482	1775	3923
GUDHI	381	6	177	1518	442	4590
RIPSER	2	1	2	5	3	1517

Table 4: Memory usage in GB for the computations summarized in Table 3. For JAVAPLEX, we indicate the value of the maximum heap size that was sufficient to perform the computation. The value that we give is an upper bound to the memory usage. For DIPHA, we indicate the maximum memory used by a single core (considering all cores). See Table 3 for details on the number of cores used.

(a) Computations on cluster

Data set	eleg	Klein	HIV	drag 2	random	genome
Size of complex	4.4×10^6	1.1×10^7	2.1×10^8	1.3×10^9	3.1×10^9	4.5×10^8
Max. dim.	2	2	2	2	8	2
JAVAPLEX (st)	< 5	< 15	> 64	> 64	> 64	> 64
DIONYSUS (st)	1.3	11.6	-	-	-	-
DIPHA (st)	0.1	0.2	2.7	4.9	-	4.8
PERSEUS	5.1	12.7	-	-	-	-
DIONYSUS (d)	0.5	1.1	-	-	-	-
DIPHA (d)	0.1	0.2	1.8	13.8	9.6	6.3
GUDHI	0.2	0.5	8.5	62.8	-	21.5
RIPSER	0.007	0.02	0.06	0.2	24.7	0.07

(b) Computations on shared-memory system

Data set	eleg	Klein	HIV	drag 2	genome	fract r
Size of complex	3.2×10^8	1.1×10^7	2.1×10^8	1.3×10^9	4.5×10^8	2.8×10^9
Max. dim.	3	2	2	2	2	3
JAVAPLEX (st)	< 600	< 15	< 700	> 700	< 700	> 700
PERSEUS	-	11.7	-	-	-	-
DIONYSUS (d)	-	1.1	16.8	134.2	-	268.5
DIPHA (d)	31.2	0.9	17.7	109.5	37.3	276.1
GUDHI	15.4	0.5	10.2	62.8	21.4	134.8
RIPSER	0.2	0.03	0.07	0.2	0.7	155

Table 5: For each software package in (a), we indicate in (b) the maximal size of the simplicial complex supported by it thus far in our tests.

(a)	JAVAPLEX	DIONYSUS		DIPHA		PERSEUS	GUDHI	RIPSER
	st	st	d	st	d	st	d	d
(b)	$4.5 \cdot 10^8$	$1.1 \cdot 10^7$	2.8×10^9	$1.3 \cdot 10^9$	$3.4 \cdot 10^9$	$1 \cdot 10^7$	$3.4 \cdot 10^9$	$3.4 \cdot 10^9$

downloaded and are then “plug-and-play,” and they have user-friendly interfaces.) Because the library JAVAPLEX implements the computation of a variety of complexes and algorithms, we feel that it is the best software for an initial foray into PH computation.

We now give guidelines for the computation of PH based on our benchmarking. We list several types of data sets in Table 6 and indicate which software or algorithm that we feel is best-suited to each one. These guidelines are based on the findings of our benchmarking. Note that one can transform networks into distance matrices, and distance matrices can yield points in Euclidean space using a method such as multi-dimensional scaling. Naturally, given a finite set of points in Euclidean space, we can compute their distance matrix. As we discussed in Section 5.1, image data can also be considered as a finite metric space, so the indications for distance matrices and points in Euclidean space also apply to image data.

Table 6: Guidelines for which implementation is best-suited for which data set, based on our benchmarking. Recall that we indicate the implementation of the dual algorithm using the abbreviation “d” following the name of a package, and similarly we indicate the implementation of the standard algorithm by “st”. Note that for smaller data sets one can also use JAVAPLEX to compute PH with VR complexes from points in Euclidean space, and PERSEUS to compute PH with cubical complexes for image data, and with VR complexes for distance matrices. The library JHOLES can only handle networks with density much less than 1.

Data type	Complexes	Suggested software
networks	WRCF	JHOLES
image data	cubical	GUDHI or DIPHA (st)
distance matrix	VR	RIPSER
distance matrix	W	JAVAPLEX
points in Euclidean space	VR	GUDHI
points in Euclidean space	\tilde{C}	DIONYSUS
points in Euclidean space	α (only in dim 2 and 3)	DIONYSUS ((st) in dim 2, (d) in dim 3) or GUDHI

8 Future directions

We conclude by discussing some future directions for the computation of PH. As we saw in Section 5, much work has been done on step 2 (i.e., going from filtered complexes to barcodes) of the PH pipeline of Fig. 6, and there exist implementations of many fast algorithms for the reduction of the boundary matrix. Step 1 (i.e., going from data to a filtered complex) of the PH pipeline is an active area of research, but many sparsification techniques (see, e.g., [79, 113]) for complexes have yet to be implemented, and more research needs to be done on steps 1 and 3 (i.e., interpreting barcodes; see, e.g., [17, 20, 122]) of the PH pipeline.

We believe that there needs to be a community-wide effort to build a library that implements the algorithms and data structures for the computation of PH, and that it should be done in a way that new algorithms and methods can be implemented easily in this framework. This would parallel similar community-wide efforts in fields such as computational algebra and computational geometry, and libraries such as Macaulay2 [66], Sage [42], and CGAL [120].

We also believe that there is a need to create guidelines and benchmark data sets for the test of new algorithms and data structures. The methods and collection of data sets that we used in our benchmarking provide an initial step towards establishing such guidelines and a list of test problems.

Acknowledgements

We thank the Rabadan Lab at Columbia University for providing the HIV and H3N2 sequences used in [26] and Giovanni Petri for sharing the data sets used in [108]. We thank Adrian Clough, Patrizio Frosini, Florian Klimm, Vitaliy Kurlin, Robert MacKay, James Meiss, Dane Taylor, Leo Speidel, Parker Edwards, and Bernadette Stolz for helpful comments on a draft of this paper. We also thank the anonymous referees for their many helpful comments. The first author thanks Ulrich Bauer, Michael Lesnick, Hubert Wagner, and Matthew Wright for helpful discussions, and thanks Florian Klimm, Vidit Nanda, and Bernadette Stolz for precious advice. The authors would like to acknowledge the use of the University of Oxford Advanced Research Computing (ARC) facility (<http://dx.doi.org/10.5281/zenodo.22558>) in carrying out some of the computations performed in this work. The first author thanks the support team at the ARC for their assistance. NO and PG are grateful for support from the EPSRC grant EP/G065802/1 (The Digital Economy HORIZON Hub). HAH gratefully acknowledges EPSRC Fellowship EP/K041096/1. NO and UT were supported by The Alan Turing Institute through EPSRC grant EP/N510129/1. The authors declares that they have no competing interests.

References

- [1] *Volvis repository*. <http://volvis.org>.
- [2] H. ADAMS AND A. TAUSZ, *JavaPlex tutorial*. Available at <https://github.com/appliedtopology/javaplex>.
- [3] A. ADCOCK, E. CARLSSON, AND G. CARLSSON, *The ring of algebraic functions on persistence bar codes*, ArXiv e-prints, (2013). 1304.0530.
- [4] D. ATTALI, H. EDELSBRUNNER, AND Y. MILEYKO, *Weak witnesses for Delaunay triangulations of submanifolds*, in Proceedings of the 2007 ACM Symposium on Solid and Physical Modeling, SPM '07, New York, NY, USA, 2007, ACM, pp. 143–150.
- [5] U. BAUER, *Ripser*, 2016. Available at <https://github.com/Ripser/ripser>.
- [6] U. BAUER, M. KERBER, AND J. REININGHAUS, *DIPHA (A distributed persistent homology algorithm)*. Software available at <https://code.google.com/p/dipha/>.
- [7] ———, *Clear and compress: Computing persistent homology in chunks*, in Topological Methods in Data Analysis and Visualization III, P.-T. Bremer, I. Hotz, V. Pascucci, and R. Peikert, eds., Mathematics and Visualization, Springer International Publishing, 2014, pp. 103–117.
- [8] ———, *Distributed computation of persistent homology*, in 2014 Proceedings of the Sixteenth Workshop on Algorithm Engineering and Experiments (ALENEX), Society for industrial and applied mathematics, 2014, ch. 3, pp. 31–38.
- [9] U. BAUER, M. KERBER, J. REININGHAUS, AND H. WAGNER, *PHAT: Persistent homology algorithms toolbox*, in Mathematical Software, ICMS 2014, H. Hong and C. Yap, eds., vol. 8592 of Lecture Notes in Computer Science, Springer Berlin Heidelberg, 2014, pp. 137–143. Software available at <https://code.google.com/p/phat/>.
- [10] P. BENDICH, H. EDELSBRUNNER, AND M. KERBER, *Computing robustness and persistence for images*, IEEE Transactions on Visualization and Computer Graphics, 16 (2010), pp. 1251–1260.

- [11] S. BHATTACHARYA, R. GHRIST, AND V. KUMAR, *Persistent homology for path planning in uncertain environments*, IEEE Transactions on Robotics, 31 (2015), pp. 578–590.
- [12] J. BINCHI, E. MERELLI, M. RUCCO, G. PETRI, AND F. VACCARINO, *jHoles: A tool for understanding biological complex networks via clique weight rank persistent homology*, Electronic Notes in Theoretical Computer Science, 306 (2014), pp. 5–18. Proceedings of the 5th International Workshop on Interactions between Computer Science and Biology (CS2Bio14).
- [13] A. BJÖRNER, *Topological methods*, in Handbook of combinatorics, R. Graham, M. Grötschel, and L. Lovász, eds., Elsevier Science B.V., 1995, ch. 34, pp. 1819–1872.
- [14] J.-D. BOISSONNAT, O. DEVILLERS, AND S. HORNUS, *Incremental construction of the Delaunay triangulation and the Delaunay graph in medium dimension*, in Proceedings of the Twenty-fifth Annual Symposium on Computational Geometry, SCG '09, New York, NY, USA, 2009, ACM, pp. 208–216.
- [15] J.-D. BOISSONNAT, L. J. GUIBAS, AND S. Y. OUDOT, *Manifold reconstruction in arbitrary dimensions using witness complexes*, Discrete & Computational Geometry, 42 (2009), pp. 37–70.
- [16] J.-D. BOISSONNAT AND C. MARIA, *Computing persistent homology with various coefficient fields in a single pass*, in Algorithms - ESA 2014, A. S. Schulz and D. Wagner, eds., vol. 8737 of Lecture Notes in Computer Science, Springer Berlin Heidelberg, 2014, pp. 185–196.
- [17] P. BUBENIK, *Statistical topological data analysis using persistence landscapes*, Journal of Machine Learning Research, 16 (2015), pp. 77–102.
- [18] P. BUBENIK, V. DE SILVA, AND J. SCOTT, *Metrics for generalized persistence modules*, Foundations of Computational Mathematics, 15 (2014), pp. 1501–1531.
- [19] P. BUBENIK AND P. DŁOTKO, *A persistence landscapes toolbox for topological statistics*, ArXiv e-prints, (2015). 1501.00179.
- [20] P. BUBENIK AND P. T. KIM, *A statistical approach to persistent homology*, Homology, Homotopy and Applications, 9 (2007), pp. 337–362.
- [21] P. BUBENIK AND J. A. SCOTT, *Categorification of persistent homology*, Discrete & Computational Geometry, 51 (2014), pp. 600–627.
- [22] G. CARLSSON, *Topology and data*, Bulletin of the American Mathematical Society, 46 (2009), pp. 255–308.
- [23] G. CARLSSON, V. DE SILVA, AND D. MOROZOV, *Zigzag persistent homology and real-valued functions*, in Proceedings of the Twenty-fifth Annual Symposium on Computational Geometry, SCG '09, New York, NY, USA, 2009, ACM, pp. 247–256.
- [24] G. CARLSSON, T. ISHKHANOV, V. DE SILVA, AND A. ZOMORODIAN, *On the local behavior of spaces of natural images*, International Journal of Computer Vision, 76 (2008), pp. 1–12.
- [25] G. CARLSSON AND A. ZOMORODIAN, *The theory of multidimensional persistence*, Discrete & Computational Geometry, 42 (2009), pp. 71–93.

- [26] J. M. CHAN, G. CARLSSON, AND R. RABADAN, *Topology of viral evolution*, Proceedings of the National Academy of Sciences of the United States of America, 110 (2013), pp. 18566–18571.
- [27] F. CHAZAL, D. COHEN-STEINER, M. GLISSE, L. J. GUIBAS, AND S. Y. OUDOT, *Proximity of persistence modules and their diagrams*, in Proceedings of the Twenty-fifth Annual Symposium on Computational Geometry, SCG '09, New York, NY, USA, 2009, ACM, pp. 237–246.
- [28] C. CHEN AND M. KERBER, *Persistent homology computation with a twist*, in Proceedings of the 27th European Workshop on Computational Geometry, 2011, pp. 197–200.
- [29] S. CHEPUSHTANOVA, T. EMERSON, E. HANSON, M. KIRBY, F. MOTTA, R. NEVILLE, C. PETERSON, P. SHIPMAN, AND L. ZIEGELMEIER, *Persistence images: An alternative persistent homology representation*, ArXiv e-prints, (2015). 1507.06217.
- [30] M. K. CHUNG, P. BUBENIK, AND P. T. KIM, *Persistence diagrams of cortical surface data*, in Information Processing in Medical Imaging, J. L. Prince, D. L. Pham, and K. J. Myers, eds., vol. 5636 of Lecture Notes in Computer Science, Springer Berlin Heidelberg, 2009, pp. 386–397.
- [31] M. M. COHEN, *A Course in Simple Homotopy Theory*, Graduate Texts in Mathematics, Springer-Verlag New York Inc., 1970.
- [32] D. COHEN-STEINER, H. EDELSBRUNNER, AND J. HARER, *Stability of persistence diagrams*, Discrete & Computational Geometry, 37 (2007), pp. 103–120.
- [33] D. COPPERSMITH AND S. WINOGRAD, *Matrix multiplication via arithmetic progressions*, Journal of Symbolic Computation, 9 (1990), pp. 251–280. Computational algebraic complexity editorial.
- [34] J. CURRY, *Sheaves, cosheaves and applications*, ArXiv e-prints, (2013). 1303.3255.
- [35] C. CURTO, *What can topology tell us about the neural code?*, Bulletin (New Series) of the American Mathematical Society, 54 (2017), pp. 63–78.
- [36] T. A. DAVIS AND Y. HU, *The University of Florida Sparse Matrix Collection*, ACM Transactions on Mathematical Software, 38 (2011), pp. 1–25. Available at <http://www.cise.ufl.edu/research/sparse/matrices>.
- [37] V. DE SILVA, *A weak characterisation of the Delaunay triangulation*, Geometriae Dedicata, 135 (2008), pp. 39–64.
- [38] V. DE SILVA AND G. CARLSSON, *Topological estimation using witness complexes*, in Proceedings of the First Eurographics conference on Point-Based Graphics, Eurographics Association, 2004, pp. 157–166.
- [39] V. DE SILVA AND R. GHRIST, *Coverage in sensor networks via persistent homology*, Algebraic & Geometric Topology, 7 (2007), pp. 339–358.
- [40] V. DE SILVA, D. MOROZOV, AND M. VEJDEMO-JOHANSSON, *Dualities in persistent (co)homology*, Inverse Problems, 27 (2011), p. 124003.
- [41] ———, *Persistent cohomology and circular coordinates*, Discrete & Computational Geometry, 45 (2011), pp. 737–759.

- [42] T. S. DEVELOPERS, *Sage Mathematics Software*. Available at <http://www.sagemath.org>.
- [43] D. DEWOSKIN, J. CLIMENT, I. CRUZ-WHITE, M. VAZQUEZ, C. PARK, AND J. ARSUAGA, *Applications of computational homology to the analysis of treatment response in breast cancer patients*, Topology and its Applications, 157 (2010), pp. 157–164. Proceedings of the International Conference on Topology and its Applications 2007 at Kyoto; Jointly with 4th Japan Mexico Topology Conference.
- [44] T. K. DEY, F. FAN, AND Y. WANG, *Graph induced complex on point data*, in Proceedings of the Twenty-ninth Annual Symposium on Computational Geometry, SoCG '13, New York, NY, USA, 2013, ACM, pp. 107–116.
- [45] ———, *Computing topological persistence for simplicial maps*, in Proceedings of the Thirtieth Annual Symposium on Computational Geometry, SOCG'14, New York, NY, USA, 2014, ACM, pp. 345–354.
- [46] T. K. DEY, D. SHI, AND Y. WANG, *SimBa: An efficient tool for approximating Rips-filtration persistence via simplicial batch-collapse*, in LIPIcs-Leibniz International Proceedings in Informatics. Vol. 57, Schloss Dagstuhl-Leibniz-Zentrum fuer Informatik, 2016.
- [47] P. DŁOTKO, *Persistence landscape toolbox*. Available at <https://www.math.upenn.edu/~dlotko/persistenceLandscape.html>.
- [48] H. EDELSBRUNNER, *The union of balls and its dual shape*, Discrete & Computational Geometry, 13 (1995), pp. 415–440.
- [49] H. EDELSBRUNNER AND M. D., *Persistent homology: Theory and practice*, in Proceedings of the European Congress of Mathematics, 2012, pp. 31–50.
- [50] H. EDELSBRUNNER AND J. HARER, *Persistent homology — A survey*, in Surveys on Discrete and Computational Geometry. Twenty Years Later, J. E. Goodman, J. Pach, and R. Pollack, eds., vol. 453 of Contemporary Mathematics, 2008, pp. 257–282.
- [51] ———, *Computational Topology: An Introduction*, Applied mathematics, American Mathematical Society, 2010.
- [52] H. EDELSBRUNNER, D. KIRKPATRICK, AND R. SEIDEL, *On the shape of a set of points in the plane*, IEEE Transactions on Information Theory, 29 (1983), pp. 551–559.
- [53] H. EDELSBRUNNER, D. LETSCHER, AND A. ZOMORODIAN, *Topological persistence and simplification*, Discrete & Computational Geometry, 28 (2002), pp. 511–533.
- [54] H. EDELSBRUNNER, D. MOROZOV, AND V. PASCUCCI, *Persistence-sensitive simplification functions on 2-manifolds*, in Proceedings of the Twenty-second Annual Symposium on Computational Geometry, SCG '06, New York, NY, USA, 2006, ACM, pp. 127–134.
- [55] H. EDELSBRUNNER AND E. P. MÜCKE, *Three-dimensional alpha shapes*, ACM Transactions on Graphics, 13 (1994), pp. 43–72.
- [56] S. EILENBERG AND N. E. STEENROD, *Foundations of Algebraic Topology*, Princeton Mathematical Series, Princeton University Press, 1952.
- [57] B. T. FASY, J. KIM, F. LECCI, AND C. MARIA, *Introduction to the R package TDA*, ArXiv e-prints, (2014). 1411.1830.

- [58] B. T. FASY, J. KIM, F. LECCI, C. MARIA, AND V. ROUVREAU, *TDA: Statistical tools for topological data analysis*. Available at <https://cran.r-project.org/web/packages/TDA/index.html>.
- [59] M. GAMEIRO, Y. HIRAOKA, S. IZUMI, M. KRAMÁR, K. MISCHAIKOW, AND V. NANDA, *A topological measurement of protein compressibility*, Japan Journal of Industrial and Applied Mathematics, 32 (2015), pp. 1–17.
- [60] G. GAN, C. MA, AND J. WU, *Data Clustering: Theory, Algorithms, and Applications*, SIAM, Philadelphia, PA, 2007.
- [61] R. GHRIST, *Barcodes: The persistent topology of data*, Bulletin of the American Mathematical Society, 45 (2008), pp. 61–75.
- [62] ———, *Elementary Applied Topology*, Createspace, 2014. ed. 1.0.
- [63] C. GIUSTI, R. GHRIST, AND D. S. BASSETT, *Two’s company and three (or more) is a simplex*, Journal of Computational Neuroscience, 41 (2016), pp. 1–14.
- [64] A. GOLDENBERG, A. X. ZHENG, S. E. FIENBERG, AND E. M. AIROLDI, *A survey of statistical network models*, Foundations and Trends in Machine Learning, 2 (2010), pp. 129–233.
- [65] J. E. GOODMAN AND J. O’ROURKE, eds., *Handbook of Discrete and Computational Geometry*, CRC Press, Inc., Boca Raton, FL, USA, 2 ed., 2004.
- [66] D. R. GRAYSON AND M. E. STILLMAN, *Macaulay2, a software system for research in algebraic geometry*. Available at <http://www.math.uiuc.edu/Macaulay2/>.
- [67] L. J. GUIBAS AND S. Y. OUDOT, *Reconstruction using witness complexes*, Discrete & Computational Geometry, 40 (2008), pp. 325–356.
- [68] M. GUILLEMARD, H. BOCHE, G. KUTYNIOK, AND F. PHILIPP, *Signal analysis with frame theory and persistent homology*, in 10th International Conference on Sampling Theory and Applications, 2013, pp. 309–312.
- [69] A. HATCHER, *Algebraic Topology*, Cambridge University Press, Cambridge, New York, 2002.
- [70] Y. HIRAOKA, A. HIRATA, E. G. ESCOLAR, AND Y. NISHIURA, *Persistent homology and many-body atomic structure for medium-range order in the glass*, Nanotechnology, 26 (2015), p. 304001.
- [71] D. HORAK, S. MALETIĆ, AND M. RAJKOVIĆ, *Persistent homology of complex networks*, Journal of Statistical Mechanics: Theory and Experiment, 2009 (2009), p. P03034.
- [72] J. JONSSON, *Simplicial Complexes of Graphs*, Lecture Notes in Mathematics, Springer Berlin Heidelberg, 2007.
- [73] M. JOSWIG AND M. E. PFETSCH, *Computing optimal Morse matchings*, SIAM Journal on Discrete Mathematics, 20 (2006), pp. 11–25.
- [74] D. O. C. S. JYAMITI RESEARCH GROUP (PROF. TAMAL K. DEY) AND O. S. U. ENGINEERING, *GIComplex*, 2013. Available at <http://web.cse.ohio-state.edu/~tamaldey/GIC/GICsoftware/>.

- [75] D. O. C. S. JYAMITI RESEARCH GROUP (PROF. TAMAL K. DEY) AND O. S. U. ENGINEERING, *SimpPers*, 2014. Available at <http://web.cse.ohio-state.edu/~tamaldey/SimpPers/SimpPers-software/>.
- [76] T. KACZYNSKI, K. MISCHAIKOW, AND M. MROZEK, *Computational Homology*, vol. 157 of Applied Mathematical Sciences, Springer, 2004.
- [77] M. KAHLE, *Random geometric complexes*, Discrete & Computational Geometry, 45 (2011), pp. 553–573.
- [78] L. KAUFMAN AND P. J. ROUSSEEUW, *Finding Groups in Data: An Introduction to Cluster Analysis*, John Wiley and Sons, Inc., New York, NY, 1990.
- [79] M. KERBER AND R. SHARATHKUMAR, *Approximate Čech complex in low and high dimensions*, in 24th International Symposium on Algorithms and Computation (ISAAC 2013), L. Cai, S.-W. Cheng, and T.-W. Lam, eds., Lecture Notes in Computer Science 8283, 2013, pp. 666–676.
- [80] V. KOVACEV-NIKOLIC, P. BUBENIK, D. NIKOLIĆ, AND G. HEO, *Using persistent homology and dynamical distances to analyze protein binding*, ArXiv e-prints, (2014). 1412.1394.
- [81] M. KRAMÁR, A. GOULLET, L. KONDIC, AND K. MISCHAIKOW, *Persistence of force networks in compressed granular media*, Physical Review E, 87 (2013), p. 042207.
- [82] ———, *Quantifying force networks in particulate systems*, Physica D, 283 (2014), pp. 37–55.
- [83] V. KURLIN, 2015. <http://kurlin.org/projects/persistent-skeletons.cpp>.
- [84] ———, *A one-dimensional homologically persistent skeleton of an unstructured point cloud in any metric space*, Computer Graphics Forum, 34 (2015), pp. 253–262.
- [85] L. A. N. LABORATORY, *HIV Database*. Available at <http://www.hiv.lanl.gov/content/index>.
- [86] S. U. C. G. LABORATORY, *The stanford 3D scanning repository*. Available at <https://graphics.stanford.edu/data/3Dscanrep>.
- [87] G. LEIBON, S. PAULS, D. ROCKMORE, AND R. SAVELL, *Topological structures in the equities market network*, Proceedings of the National Academy of Sciences of the United States of America, 105 (2008), pp. 20589–20594.
- [88] M. LESNICK AND M. WRIGHT, *Interactive visualization of 2-D persistence modules*, ArXiv e-prints, (2015). 1512.00180.
- [89] M. LESNICK AND M. WRIGHT, *RIVET: The rank invariant visualization and exploration tool*, 2016. Software available at <http://rivet.online/>.
- [90] D. LO AND B. PARK, *Modeling the spread of the Zika virus using topological data analysis*, ArXiv e-prints, (2016). 1612.03554.
- [91] S. MALETIC, Y. ZHAO, AND M. RAJKOVIC, *Persistent topological features of dynamical systems*, ArXiv e-prints, (2015). 1510.06933.
- [92] C. MARIA, *Algorithms and data structures in computational topology*, (2014). PhD Thesis, Université de Nice-Sophia Antipolis. Available at http://www-sop.inria.fr/members/Clement.Maria/docs/ClementMaria_PhDdissertation.pdf.

- [93] C. MARIA, J.-D. BOISSONNAT, M. GLISSE, AND M. YVINEC, *The Gudhi library: Simplicial complexes and persistent homology*, in Mathematical Software, ICMS 2014, H. Hong and C. Yap, eds., vol. 8592 of Lecture Notes in Computer Science, Springer Berlin Heidelberg, 2014, pp. 167–174. Software available at <https://project.inria.fr/gudhi/software/>.
- [94] Y. MILEYKO, S. MUKHERJEE, AND J. HARER, *Probability measures on the space of persistence diagrams*, Inverse Problems, 27 (2011), p. 124007.
- [95] N. MILOSAVLJEVIĆ, D. MOROZOV, AND P. SKRABA, *Zigzag persistent homology in matrix multiplication time*, in Proceedings of the Twenty-seventh Annual Symposium on Computational Geometry, SoCG '11, New York, NY, USA, 2011, ACM, pp. 216–225.
- [96] K. MISCHAIKOW AND V. NANDA, *Morse theory for filtrations and efficient computation of persistent homology*, Discrete & Computational Geometry, 50 (2013), pp. 330–353.
- [97] D. MOROZOV, *Dionysus*. Software available at <http://www.mrzv.org/software/dionysus/>.
- [98] ———, *Persistence algorithm takes cubic time in worst case*, BioGeometry News, Dept. Comput. Sci., Duke Univ, (2005).
- [99] E. MUNCH, K. TURNER, P. BENDICH, S. MUKHERJEE, J. MATTINGLY, AND J. HARER, *Probabilistic Fréchet means for time varying persistence diagrams*, Electronic Journal of Statistics, 9 (2015), pp. 1173–1204.
- [100] V. NANDA, *Perseus, the persistent homology software*. Software available at <http://www.sas.upenn.edu/~vnanda/perseus>.
- [101] ———, *Discrete morse theory for filtrations*, (2012). PhD Thesis. Rutgers, The State University of New Jersey.
- [102] M. E. J. NEWMAN, *Finding community structure in networks using the eigenvectors of matrices*, Physical Review E, 74 (2006), p. 036104.
- [103] M. NICOLAU, A. J. LEVINE, AND G. CARLSSON, *Topology based data analysis identifies a subgroup of breast cancers with a unique mutational profile and excellent survival*, Proceedings of the National Academy of Sciences of the United States of America, 108 (2011), pp. 7265–7270.
- [104] S. Y. OUDOT, *Persistence Theory: From Quiver Representations to Data Analysis*, vol. 209 of AMS Mathematical Surveys and Monographs, American Mathematical Society, 2015.
- [105] M. PENROSE, *Random Geometric Graphs*, Oxford University Press, Oxford, UK, 2003.
- [106] J. A. PEREA, A. DECKARD, S. B. HAASE, AND J. HARER, *Sw1pers: Sliding windows and 1-persistence scoring; discovering periodicity in gene expression time series data*, BMC Bioinformatics, 16 (2015), pp. 1–12.
- [107] P. PERRY AND V. DE SILVA, *Plex*, 2000–2006. Available at <http://mii.stanford.edu/research/comptop/programs/>.
- [108] G. PETRI, M. SCOLAMIERO, I. DONATO, AND F. VACCARINO, *Topological strata of weighted complex networks*, PLoS ONE, 8 (2013), pp. 1–8.

- [109] F. T. POKORNY, M. HAWASLY, AND S. RAMAMOORTHY, *Topological trajectory classification with filtrations of simplicial complexes and persistent homology*, The International Journal of Robotics Research, 35 (2016), pp. 204–223.
- [110] K. T. POOLE, *Voteview*. <http://voteview.com>, 2016.
- [111] F. ROBIN, *Morse theory for cell complexes*, Advances in Mathematics, 134 (1998), pp. 90–145.
- [112] S. E. SCHAEFFER, *Graph clustering*, Computer Science Review, 1 (2007), pp. 27–64.
- [113] D. R. SHEEHY, *Linear-size approximations to the Vietoris-Rips filtration*, Discrete & Computational Geometry, 49 (2013), pp. 778–796.
- [114] G. SINGH, F. MÉMOLI, AND G. CARLSSON, *Topological methods for the analysis of high dimensional data sets and 3D object recognition*, in Eurographics Symposium on Point-Based Graphics, 2007, pp. 91–100.
- [115] N. SINGH, H. D. COUTURE, J. S. MARRON, C. PEROU, AND M. NIETHAMMER, *Topological descriptors of histology images*, in Machine Learning in Medical Imaging, G. Wu, D. Zhang, and L. Zhou, eds., vol. 8679 of Lecture Notes in Computer Science, Springer International Publishing, 2014, pp. 231–239.
- [116] O. SPORNS, *Small-world connectivity, motif composition, and complexity of fractal neuronal connections*, Biosystems, 85 (2006), pp. 55–64.
- [117] B. J. STOLZ, H. A. HARRINGTON, AND M. A. PORTER, *Persistent homology of time-dependent functional networks constructed from coupled time series*, ArXiv e-prints, (2016). 1605.00562.
- [118] A. TAUSZ, M. VEJDEMO-JOHANSSON, AND H. ADAMS, *JavaPlex: A research software package for persistent (co)homology*, in Proceedings of ICMS 2014, H. Hong and C. Yap, eds., Lecture Notes in Computer Science 8592, 2014, pp. 129–136. Software available at <http://appliedtopology.github.io/javaplex/>.
- [119] D. TAYLOR, F. KLIMM, H. A. HARRINGTON, M. KRAMÁR, K. MISCHAIKOW, M. A. PORTER, AND P. J. MUCHA, *Topological data analysis of contagion maps for examining spreading processes on networks*, Nature Communications, 6 (2015). Article number 7723.
- [120] THE CGAL PROJECT, *CGAL user and reference manual*, CGAL Editorial Board, 4.7 ed., 2015.
- [121] C. M. TOPAZ, L. ZIEGELMEIER, AND T. HALVERSON, *Topological data analysis of biological aggregation models*, PLoS ONE, 10 (2015), pp. 1–26.
- [122] K. TURNER, Y. MILEYKO, S. MUKHERJEE, AND J. HARER, *Fréchet means for distributions of persistence diagrams*, Discrete & Computational Geometry, 52 (2014), pp. 44–70.
- [123] R. VASUDEVAN, A. AMES, AND R. BAJCSY, *Persistent homology for automatic determination of human-data based cost of bipedal walking*, Nonlinear Analysis: Hybrid Systems, 7 (2013), pp. 101–115. IFAC World Congress 2011.
- [124] T. VICSEK, A. CZIRÓK, E. BEN-JACOB, I. COHEN, AND O. SHOCHET, *Novel type of phase transition in a system of self-driven particles*, Physical Review Letters, 75 (1995), pp. 1226–1229.

- [125] L. VIETORIS, *Über den höheren Zusammenhang kompakter Räume und eine Klasse von zusammenhangstreuen Abbildungen*, Mathematische Annalen, 97 (1927), pp. 454–472.
- [126] H. WAGNER, C. CHEN, AND E. VUČINI, *Efficient computation of persistent homology for cubical data*, in Topological Methods in Data Analysis and Visualization II, R. Peikert, H. Hauser, H. Carr, and R. Fuchs, eds., Mathematics and Visualization, Springer Berlin Heidelberg, 2012, pp. 91–106.
- [127] B. WANG AND G.-W. WEI, *Object-oriented persistent homology*, Journal of Computational Physics, 305 (2016), pp. 276–299.
- [128] A. S. WAUGH, L. PEI, J. H. FOWLER, P. J. MUCHA, AND M. A. PORTER, *Party polarization in congress: A network science approach*, arXiv:0907.3509, (2012). Data available at http://figshare.com/articles/Roll_Call_Votes_United_States_House_and_Senate/1590036.
- [129] K. XIA, X. FENG, Y. TONG, AND G. W. WEI, *Persistent homology for the quantitative prediction of fullerene stability*, Journal of Computational Chemistry, 36 (2015), pp. 408–422.
- [130] K. XIA AND G.-W. WEI, *Persistent homology analysis of protein structure, flexibility, and folding*, International Journal for Numerical Methods in Biomedical Engineering, 30 (2014), pp. 814–844.
- [131] K. XIA AND G.-W. WEI, *A review of geometric, topological and graph theory apparatuses for the modeling and analysis of biomolecular data*, ArXiv e-prints, (2016). 1612.01735.
- [132] W. ZHOU AND H. YAN, *Alpha shape and Delaunay triangulation in studies of protein-related interactions*, Briefings in Bioinformatics, (2012).
- [133] A. ZOMORODIAN, *Topology for Computing*, Cambridge Monographs on Applied and Computational Mathematics, Cambridge University Press, 2009.
- [134] —, *Technical section: Fast construction of the Vietoris-Rips complex*, Comput. Graph., 34 (2010), pp. 263–271.
- [135] —, *The tidy set: A minimal simplicial set for computing homology of clique complexes*, 2010, pp. 257–266. SCG10.
- [136] —, *Topological data analysis*, in Advances in Applied and Computational Topology: American Mathematical Society Short Course on Computational Topology, January 4-5, 2011, New Orleans, Louisiana, A. Zomorodian, ed., vol. 70 of Proceedings of symposia in applied mathematics, American Mathematical Society, 2012, pp. 1–39.
- [137] A. ZOMORODIAN AND G. CARLSSON, *Computing persistent homology*, Discrete Comput. Geom., 33 (2005), pp. 249–274.